

POLITECNICO DI TORINO

Master's Degree in Energy and Nuclear Engineering



**Politecnico
di Torino**

Non-perturbative computation of neutron noise using the Monte Carlo method

In collaboration with:



Supervisors

Prof. Sandra DULLA

Mr. Axel FAUVEL

Dr. Davide MANCUSI

Dr. Amélie ROUCHON

Dr. Andrea ZOIA

Candidate

Olivia FERRERO

a.y. 2025/2026

Abstract

Neutron noise denotes small time-dependent fluctuations of the neutron flux around its stationary value, typically induced by perturbations in the reactor core. These perturbations may originate from changes in coolant density, vibrations of fuel assemblies or control rods, and other mechanical or operational effects. In power reactors, neutron noise can be monitored through in-core and ex-core detectors for diagnostics, anomaly detection, and operational surveillance. This thesis investigates the non-perturbative Monte Carlo formulation of neutron noise in the frequency domain and benchmarks it against the standard, or orthodox, linearised approach. All simulations were performed with the Monte Carlo mini-app MGMC, developed at CEA, which supports both linear and non-perturbative neutron noise calculations. The analysis first considers simplified configurations based on infinite homogeneous media, in order to compare the two formulations under controlled conditions and to assess the role of harmonic coupling. Increasing levels of complexity are then introduced by considering multiple energy groups and perturbations with different harmonic content. The study is subsequently extended to a one-dimensional vibrating-rod model, representative of the noise induced by mechanical vibrations. For this benchmark, a reformulated non-perturbative eigenvalue problem is adopted in order to make the Monte Carlo simulations viable for realistic vibration amplitudes. The results show that the linear and non-perturbative formulations remain in close agreement for the dominant components of the response, whereas the non-perturbative treatment becomes necessary when harmonic coupling and higher-order harmonics play a significant role. Comparisons with deterministic reference solutions confirm the accuracy of the non-perturbative Monte Carlo methodology for the dominant components of the response, while also highlighting the remaining difficulty of resolving weak higher-order harmonics.

Acknowledgements

First of all, I would like to express my deepest gratitude to my supervisor, Axel, for his guidance, support, and trust throughout this work. I am especially thankful for his availability and patience. His advice and feedback were invaluable, and working under his supervision has been a great opportunity for growth.

I am also sincerely grateful to my other supervisors, Andrea, Davide, and Amélie, for their valuable advice, availability, and support throughout this experience.

My sincere gratitude also goes to Professor Dulla, my thesis advisor, for her support and availability throughout this path. I am especially thankful to her for being such an inspiration model.

I would also like to warmly thank all the people I met at CEA, not only for creating such a stimulating working environment, but also for all the moments we shared outside of work, which made this period especially meaningful and enjoyable.

A special thank you goes to my roommates, Luigi and Enrico, with whom I had the chance to share both everyday life and this experience at CEA. Their presence made this journey lighter and even more memorable. I am also very grateful to Thomas and Bulga for all the time we spent together, for the work we shared, and for the many enjoyable moments along the way.

My heartfelt thanks go to my friends Giuseppe and Andrea, with whom I shared many important moments throughout my years at Politecnico. I am grateful for their friendship, support, and for all the experiences we went through together.

I am deeply thankful to my parents and my sister Vanessa for their constant love and support, and for always being there for me.

Finally, I would like to thank my boyfriend, Simone, for always being by my side, for his constant support and patience, and for helping me through all the difficult moments of this journey.

Table of Contents

List of Tables	IX
List of Figures	XI
1 Introduction	1
1.1 Context and objectives	1
1.2 Organisation of the thesis	3
2 Neutron Transport Theory	5
2.1 Nuclear fission reactors	5
2.2 Nuclear fission	6
2.2.1 Chain reaction and criticality	7
2.3 Transport theory	8
2.3.1 Neutron interactions	9
2.3.2 Phase space	10
2.3.3 Neutron cross sections	11
2.3.4 Fission source parameters	12
2.3.5 Neutron flux and current	13
2.4 Neutron transport equation	13
2.4.1 Operator form of the transport equation	14
2.4.2 Steady-state transport equation	15

2.5	Multi-group approximation	15
2.5.1	Useful definitions	16
2.5.2	Multi-group transport equation	17
2.6	Kinetic approximations	18
3	Monte Carlo Method	21
3.1	Statistical foundations	22
3.1.1	Random variables, distributions and moments	22
3.1.2	Sample mean and statistical uncertainty	23
3.1.3	Sample variance	24
3.1.4	Bias and figure of merit	24
3.2	Monte Carlo solution of the transport equation	25
3.2.1	Particle histories and random walks	25
3.2.2	Sampling of free flight and collision type	25
3.2.3	Analog and non-analog transport	26
3.3	Estimators in neutron transport	27
3.3.1	General notion of a tally	27
3.3.2	Track-length estimator	27
3.3.3	Collision estimator	28
3.4	Monte Carlo eigenvalue calculations	28
3.4.1	The k -eigenvalue problem	28
3.4.2	Inactive generations, source convergence and population control	29
4	Neutron Noise Theory	31
4.1	Origins of neutron noise	31
4.2	Power reactor noise	32
4.2.1	Linear theory	33
4.2.2	Non-perturbative theory	34

4.2.3	Point-kinetics interpretation and zero power transfer function	37
5	Monte Carlo for Neutron Noise in MGMC	39
5.1	Frequency-domain Monte Carlo formulation in MGMC	39
5.1.1	Treatment of the $i\omega/v$ term	40
5.1.2	Flux estimators	41
5.1.3	Linear noise mode	42
5.1.4	Non-perturbative noise mode	42
5.1.5	Symmetry and frequency mirror	44
5.1.6	Weight cancellation	45
6	Homogeneous Medium	47
6.1	Verification of the flight-based scheme with track-length estimator .	47
6.1.1	Scattering-free infinite-medium test problem	48
6.1.2	Results before and after the correction	49
6.2	Linear vs non-perturbative comparison	51
6.2.1	Problem definition and input data	51
6.2.2	Numerical computation settings	53
6.2.3	One-group, single-harmonic perturbation	54
6.2.4	One-group, two-harmonic perturbation	57
6.2.5	Two-group, single-harmonic perturbation	59
7	One-dimensional Vibrating Rod	63
7.1	Rod model: problem definition and numerical set up	64
7.2	Numerical computation settings	66
7.3	Results for the vibrating rod benchmark	67
7.3.1	Zero-frequency component	67
7.3.2	First harmonic	68
7.3.3	Second harmonic	68

7.3.4	Third harmonic	69
7.3.5	Summary of the comparison	70
7.4	Parametric study	71
7.4.1	Inactive generations: monitoring k' and harmonic entropy	71
7.4.2	Harmonic truncation: choice of $h_{\max}$	72
7.4.3	Figure of Merit study	75
8	Conclusions and Perspectives	77
A	Mathematical Background	81
A.1	Fourier transform	81
A.1.1	Remarks on function spaces	81
A.1.2	Basic properties	82
A.1.3	Fourier transform of a periodic function	83
A.2	Convolution product	83
A.2.1	Continuous convolution	83
A.2.2	Discrete convolution	84
	Bibliography	85

List of Tables

6.1	Macroscopic cross sections and kinetic parameters for the one-group, scattering-free validation case.	49
6.2	Scattering-free test at $f = 10^4$ Hz: comparison between analytical reference and Monte Carlo fluxes before correcting the implementation of the track-length estimator.	50
6.3	Scattering-free test at $f = 10^4$ Hz: analytical reference and corrected track-length estimator.	50
6.4	Macroscopic cross sections and kinetic parameters for the one-group homogeneous configuration.	52
6.5	Macroscopic cross sections, scattering matrix, neutron speeds, and kinetic parameters for the two-group homogeneous configuration. . .	52
7.1	Input data for the one-dimensional vibrating rod model.	65
7.2	Figure of Merit ratios for the vibrating rod model.	75

List of Figures

2.1	Mass distribution of fission products from thermal fission of ^{235}U [17].	7
2.2	Neutron-induced cross sections of ^{238}U [23].	9
2.3	Position and direction of flight of a neutron [15].	10
2.4	Energy group discretisation [15].	16
4.1	Amplitude and phase of the zero power reactor transfer function $G_0(\omega)$ [1].	38
6.1	One-group, single-harmonic case: amplitude (left) and phase (right) of the normalised first harmonic as a function of frequency.	55
6.2	One-group, single-harmonic case: modulus (left) and phase (right) of $\hat{\varphi}_h$ as a function of harmonic index h at a representative plateau frequency ($f = 10$ Hz).	56
6.3	One-group, single-harmonic case: modulus (left) and phase (right) of $\hat{\varphi}_h$ as a function of harmonic index h at a representative high frequency ($f = 10^3$ Hz).	56
6.4	One-group, two-harmonic case: amplitude (left) and phase (right) of the first harmonic $\hat{\varphi}_1$ as a function of frequency.	57
6.5	One-group, two-harmonic case: amplitude (left) and phase (right) of the second harmonic $\hat{\varphi}_2$ as a function of frequency.	58
6.6	One-group, two-harmonic case: modulus (left) and phase (right) of $\hat{\varphi}_h$ as a function of harmonic index h at a representative plateau frequency ($f = 10$ Hz).	58

6.7	One-group, two-harmonic case: modulus (left) and phase (right) of $\hat{\varphi}_h$ as a function of harmonic index h at a representative high frequency ($f = 10^3$ Hz).	59
6.8	Two-group, single-harmonic case: amplitude (left) and phase (right) of the first-harmonic fast-group flux $\hat{\varphi}_{1,0}$ as a function of frequency.	60
6.9	Two-group, single-harmonic case: amplitude (left) and phase (right) of the first-harmonic thermal-group flux $\hat{\varphi}_{1,1}$ as a function of frequency.	60
6.10	Two-group, single-harmonic case: modulus (left) and phase (right) of the thermal-group coefficients $\hat{\varphi}_{h,1}$ as a function of harmonic index h at a plateau frequency ($f = 10$ Hz).	61
6.11	Two-group, single-harmonic case: modulus (left) and phase (right) of the thermal-group coefficients $\hat{\varphi}_{h,1}$ as a function of harmonic index h at a high frequency ($f = 10^2$ Hz).	61
7.1	Schematic representation of the one-dimensional vibrating rod model.	65
7.2	Real part of the zero frequency component $\hat{\varphi}_0(x)$ (top) and corresponding difference $\Re[\Delta\hat{\varphi}_0(x)]$ between Monte Carlo and deterministic solutions (bottom).	67
7.3	First harmonic $\hat{\varphi}_1(x)$ of the vibrating rod model: real part (left) and imaginary part (right). Each panel shows the normalised harmonic (top) and the difference between Monte Carlo and deterministic solutions (bottom).	68
7.4	Second harmonic $\hat{\varphi}_2(x)$ of the vibrating rod model: real part (left) and imaginary part (right). Each panel shows the normalised harmonic (top) and the difference between Monte Carlo and deterministic solutions (bottom).	69
7.5	Third harmonic $\hat{\varphi}_3(x)$ of the vibrating rod model: real part (left) and imaginary part (right). Each panel shows the normalised harmonic (top) and the difference between Monte Carlo and deterministic solutions (bottom).	70
7.6	Convergence of the non-perturbative eigenvalue k' and the harmonic entropy H as functions of generation.	72
7.7	Effect of harmonic truncation on the non-perturbative solution for $h = 1$. Relative difference $ \Delta\hat{\varphi}_1/\hat{\varphi}_1 $ with respect to a reference non-perturbative calculation with $h_{\max} = 7$, as defined in Eq. (7.3).	73

- 7.8 Effect of harmonic truncation on the non-perturbative solution for $h = 2$. Relative difference $|\Delta\hat{\varphi}_2/\hat{\varphi}_2|$ with respect to a reference non-perturbative calculation with $h_{\text{max}} = 7$, as defined in Eq. (7.3). 73
- 7.9 Effect of harmonic truncation on the non-perturbative solution for $h = 3$. Relative difference $|\Delta\hat{\varphi}_3/\hat{\varphi}_3|$ with respect to a reference non-perturbative calculation with $h_{\text{max}} = 7$, as defined in Eq. (7.3). 74

Chapter 1

Introduction

1.1 Context and objectives

Neutron noise denotes small fluctuations of the neutron flux around its steady-state value, induced by time-dependent changes in material properties, geometry, or boundary conditions [1, 2, 3]. In power reactors, these fluctuations carry useful information for core surveillance, diagnostics, and early anomaly detection. Their prediction requires the solution of the time-dependent equations for neutron transport and precursors, under periodic perturbations [1].

The pioneering works on neutron noise were based on single speed diffusion theory and on closed-form solutions obtained under strong simplifying assumptions [1, 2, 3]. More recent developments have progressively increased the physical fidelity of the models. Frequency-domain formulations were first developed within multi-group diffusion theory [4], and were later extended to transport theory [5, 6, 7]. In parallel, Monte Carlo methods were introduced in the frequency domain by using particles with complex statistical weights in a flight-based scheme [8], and were subsequently reformulated in a collision-based framework [9]. Further advances addressed more efficient sampling strategies for the evaluation of complex terms [9] and variance-reduction techniques based on weight cancellation [10].

These developments belong to the so-called ‘orthodox’, or linear, formulation of neutron noise theory, in which the underlying Fourier-transformed Boltzmann equation is simplified under the assumption of small perturbations by neglecting products of perturbed quantities [1]. This approximation has proved effective in many applications. However, it is known to become inadequate in specific situations, most notably when the noise is induced by mechanical vibrations [11, 12, 13]. In such cases, harmonic coupling plays a central role and higher harmonics may carry

essential physical information.

To address these limitations, deterministic non-perturbative formulations have been developed that retain the coupling among harmonics and provide a more consistent description of vibration-induced noise [5, 13]. These approaches make it possible to study higher harmonics explicitly and to analyse regimes in which the linear approximation is no longer sufficient. More recently, a proof of concept non-perturbative Monte Carlo strategy has been proposed for the direct solution of the neutron noise equations in the frequency domain, and has been validated on a simple homogeneous configuration [14].

The main objective of this thesis is to benchmark the non-perturbative Monte Carlo approach proposed by Fauvel *et al.* [14] against the orthodox linear formulation. To this end, the Monte Carlo mini-app MGMCMC, developed at CEA, is used to perform both linear and non-perturbative neutron noise calculations.

The assessment is carried out in two steps. First, benchmarks are performed in a zero-dimensional infinite homogeneous medium with oscillating macroscopic cross sections. This configuration provides a controlled framework for frequency-by-frequency comparisons between the two formulations and for verification against analytical solutions. Second, the analysis is extended to a one-dimensional ‘rod model’, in which the noise is induced by mechanical vibrations of a fuel rod. The adopted configuration follows the centrally vibrating fuel rod benchmark introduced by Zoia *et al.* [13]. In this case, harmonic coupling is essential, and the limitations of the linearised approach are expected to be most apparent.

For the one-dimensional study, a reformulated non-perturbative eigenvalue problem is adopted in order to improve numerical feasibility and generation control, and to enable simulations at larger vibration amplitudes. The Monte Carlo solutions obtained in this work are finally compared against an independent deterministic solver for the same model.

The original contributions of this work are fourfold. First, a dedicated scattering-free high-frequency benchmark was designed to isolate the flight-based scheme combined with the track-length estimator in MGMCMC, and to verify the corrected implementation against an analytical reference solution. Second, a systematic comparison between linear and non-perturbative Monte Carlo calculations was carried out in homogeneous one-group and two-group configurations, with deterministic reference solutions used to assess the results. Third, the reformulated non-perturbative eigenvalue formulation was applied to the one-dimensional vibrating rod benchmark, making realistic vibration amplitudes computationally accessible in Monte Carlo simulations. Finally, dedicated numerical studies were performed to analyse source convergence, harmonic truncation, and computational efficiency through Figure of

Merit estimates. Part of this work has led to a contribution to the proceedings of the PHYSOR 2026 conference.

1.2 Organisation of the thesis

This manuscript is organised as follows.

Chapter 1 presents the context, objectives, and scope of the work.

Chapter 2 recalls the essential concepts of neutron transport theory used in the thesis. The continuous-energy and multi-group Boltzmann equations are introduced, and the operator form used throughout the manuscript is defined, with particular attention to the quantities relevant to noise analysis.

Chapter 3 reviews the Monte Carlo method for particle transport and the basic probabilistic and statistical concepts that support it. The chapter discusses random walks in phase space, standard estimators for fluxes and reaction rates, and estimators of statistical uncertainty.

Chapter 4 is devoted to neutron noise theory. Both the orthodox linear formulation and the non-perturbative formulation are presented in the time and frequency domains, with emphasis on the structure of the corresponding transport equations. A reformulated non-perturbative eigenvalue problem is also introduced, which will later be used for the vibrating rod benchmark in order to overcome the convergence difficulties encountered with the original non-perturbative scheme.

Chapter 5 describes the implementation of these frequency domain formulations in the Monte Carlo mini-app MGMC. The linear and non-perturbative noise modes are detailed, including the treatment of the $i\omega/v$ term through flight-based and collision-based schemes, the handling of delayed fission, the sampling of perturbation reactions over the harmonic index, and the use of weight cancellation in the non-perturbative eigenvalue algorithm.

Chapter 6 introduces the homogeneous zero-dimensional benchmark used as a minimal test-bed for the first numerical studies. The chapter first addresses a targeted verification of the flight-based scheme combined with the track-length estimator. A dedicated scattering-free high frequency configuration is used to make a previously identified implementation error observable and to verify the corrected implementation against an analytical reference. The remainder of the chapter is then devoted to the comparison between the linear and non-perturbative formulations in one-group and two-group homogeneous configurations. The frequency response of the first harmonic and the harmonic content generated by the non-perturbative solution are also compared with deterministic reference calculations.

Chapter 7 extends the analysis to a one-dimensional vibrating rod model, representative of realistic core vibration problems. The reformulated non-perturbative eigenvalue formulation is applied to the vibrating rod configuration, and linear and non-perturbative Monte Carlo results are compared with deterministic solutions for the first harmonics of the noise field. The chapter also includes parametric studies on source convergence and harmonic truncation, together with a Figure of Merit analysis assessing the computational efficiency of the non-perturbative calculation.

Finally, Chapter 8 summarises the main achievements of the work. The capabilities and current limitations of the non-perturbative Monte Carlo implementation are discussed, and possible directions for further methodological developments and applications to more realistic reactor configurations are outlined.

Chapter 2

Neutron Transport Theory

This chapter summarises the reactor physics and neutron transport concepts that are used in the remainder of the thesis. It first recalls the physical basis of neutron multiplication in fission reactors and then introduces the transport description of neutron behaviour in phase space. The time-dependent and steady-state transport equations are presented, together with their operator form, followed by the multi-group approximation and a brief reminder of point kinetics. These elements establish the notation used in the following chapters on neutron noise theory and its Monte Carlo treatment.

2.1 Nuclear fission reactors

A nuclear fission reactor is a system designed to sustain a controlled chain reaction, thereby converting the energy released by nuclear fission into a useful form, most commonly heat and then electricity [15]. In the context of this thesis, the most relevant class is that of power reactors, since neutron noise analysis is primarily developed for reactor core monitoring and diagnostics.

From the neutronic point of view, the main distinction is between *thermal reactors* and *fast reactors*. In thermal reactors, neutrons released by fission are slowed down by a moderator, typically light water, heavy water or graphite, so that most fissions occur at low energy. In fast reactors, no moderator is used and the chain reaction is sustained by fast neutrons [16, 17].

The most widespread thermal reactor types are Pressurised Water Reactors (PWRs), Boiling Water Reactors (BWRs), Pressurised Heavy Water Reactors (PHWRs) and gas-cooled high-temperature graphite reactors (GCRs or HTGRs). PWRs and BWRs use light water as both coolant and moderator, whereas PHWRs use heavy

water and can therefore operate with natural uranium fuel. Gas-cooled reactors use graphite as moderator and a gas, typically carbon dioxide or helium, as coolant [16, 18, 19]. Fast reactors instead use coolants with weak moderating power, such as sodium or lead, in order to preserve the fast neutron spectrum [16].

Despite their technological differences, these systems share the same basic neutronic features: fissile fuel, a geometry that permits neutron multiplication, materials that scatter or absorb neutrons, a coolant, and control systems that regulate the reactivity. The common framework used to describe neutron production, loss and migration in all these systems is neutron transport theory.

2.2 Nuclear fission

Nuclear fission is a neutron-induced exothermic reaction in which a heavy nucleus absorbs a neutron, forms an excited compound nucleus, and splits into lighter fragments while releasing energy, neutrons and gamma radiation [17, 16]. In thermal reactors, the most important fissile nuclide is ^{235}U , although other relevant fissile isotopes include ^{239}Pu , ^{241}Pu and ^{233}U [16].

The average number of neutrons emitted per fission is denoted by ν . For thermal neutron induced fission of ^{235}U , ν is about 2.42 and increases moderately with incident neutron energy [16]. Most of these neutrons are emitted essentially instantaneously and are called *prompt neutrons*. A small fraction is emitted later following the radioactive decay of some fission products, and those neutrons are therefore called *delayed* [15, 16]. Denoting by ν_p and ν_d the average numbers of prompt and delayed neutrons, respectively, one has [20]

$$\nu = \nu_p + \nu_d, \quad \beta = \frac{\nu_d}{\nu}, \quad (2.1)$$

where β is the delayed neutron fraction.

Each fission of ^{235}U releases about 200 MeV of energy [16, 17]. Most of this energy appears as kinetic energy of the fission fragments, which is rapidly converted into heat in the fuel. Smaller contributions arise from prompt neutrons, prompt gamma rays and the subsequent radioactive decay of fission products. The only significant fraction that is not recoverable as heat is the energy carried away by antineutrinos [16, 17].

Neutron energy ranges. In reactor physics it is customary to distinguish three broad neutron energy ranges. *Thermal neutrons* are in approximate thermal equilibrium with the moderator and are typically associated with energies up to about 1 eV [20]. *Epithermal* or *resonance neutrons* lie above the thermal range and

are strongly affected by resonance absorption, especially in nuclides such as ^{238}U [20]. *Fast neutrons* occupy the highest energies and are described by the fission spectrum, with typical energies of the order of a few MeV [16].

It is also useful to distinguish between *fissionable*, *fissile* and *fertile* nuclides. A nuclide is *fissionable* if it can undergo fission for sufficiently energetic incident neutrons, *fissile* if it can fission with thermal neutrons, and *fertile* if it can be converted into a fissile nuclide through neutron capture followed by radioactive decay [16, 17].

The fission fragments produced in a given event are not unique. Rather, the products follow a distribution of mass yields, an example of which is shown in Fig. 2.1. The distribution is bimodal, with two peaks centred respectively around $A = 95$ and $A = 135$.




Figure 2.1: Mass distribution of fission products from thermal fission of ^{235}U [17].

2.2.1 Chain reaction and criticality

The neutrons produced at a fission event may induce further fissions. The life of the neutron population can thus be decomposed into successive fission generations. This is the basis of the nuclear chain reaction [15]. The central requirement for reactor operation is that neutron production balances neutron loss through absorption and leakage.

This balance is characterised by the effective multiplication factor k_{eff} , defined as

[15]

$$k_{\text{eff}} = \frac{\text{number of neutrons in generation } (n+1)}{\text{number of neutrons in generation } n}. \quad (2.2)$$

If $k_{\text{eff}} < 1$, the system is *subcritical*; if $k_{\text{eff}} = 1$, it is *critical*; and if $k_{\text{eff}} > 1$, it is *supercritical* [15].

A related quantity is the *reactivity* [16],

$$\rho = \frac{k_{\text{eff}} - 1}{k_{\text{eff}}}, \quad (2.3)$$

which measures the departure from criticality. A positive reactivity corresponds to a supercritical state, a negative reactivity to a subcritical state, and $\rho = 0$ to criticality [16].

Importance of delayed neutrons. Although delayed neutrons constitute only a small fraction of the total neutron population, they are essential for reactor control, because they increase the characteristic response time of the system from prompt neutron timescales to timescales accessible to mechanical control and protection systems, thanks to the ‘delay’ introduced by the presence of precursor decay [15, 16]. In normal power reactor operation, the reactor is designed and controlled so that criticality relies on delayed neutrons. When $\rho = \beta_{\text{eff}}$, the reactor becomes prompt critical. Beyond that threshold, the power rises at the prompt neutron timescale, and the transient becomes much more difficult to control [16].

2.3 Transport theory

Neutron transport theory describes the migration of neutrons and their interactions with matter through the linear Boltzmann equation. In this framework, a neutron is treated as a neutral point particle characterised by its position $\vec{r}$, direction of flight $\vec{\Omega}$, kinetic energy E , and time t [21, 22]. Between interactions, the neutron is assumed to move along a straight trajectory at constant speed. Interactions with matter are modelled as localised and instantaneous events in which the neutron may scatter, be absorbed, or induce fission [22].

The transport process is assumed to be Markovian: the probability of the next event depends only on the present state of the neutron and not on its previous history [22]. Moreover, transport observables such as flux, current and reaction rates are not assigned to individual neutrons but to an ensemble of particles. The transport equation therefore describes the evolution of ensemble-averaged quantities, while random fluctuations around these averages are neglected [21, 22]. This deterministic description provides the reference framework on which neutron noise theory is later built.

2.3.1 Neutron interactions

Neutron interactions with nuclei are governed primarily by the nuclear force. The most important reaction channels in reactor physics are the following.

Elastic scattering. In elastic scattering, the neutron changes direction and usually transfers part of its kinetic energy to the target nucleus, but the nucleus remains in its ground state. Total kinetic energy and momentum are conserved [16].

Inelastic scattering. In inelastic scattering, the neutron excites the target nucleus and therefore loses part of its kinetic energy. The excited nucleus later de-excites through gamma emission [16].

Radiative capture. In radiative capture, also denoted as an (n, γ) reaction, the neutron is absorbed and the compound nucleus emits one or more gamma rays. The neutron is removed from the neutron population [16].

Fission. In a fission reaction, neutron absorption leads to the splitting of the nucleus and to the emission of energy, neutrons and gamma radiation [15]. This process has already been discussed in Section 2.2.

Absorption with charged particle emission and threshold reactions. Neutron-induced reactions such as (n, p) and (n, α) , or $(n, 2n)$ and $(n, 3n)$ may occur. The latter two are threshold reactions and become possible only above a minimum incident neutron energy [16].

An example of the energy dependence of several reaction channels is shown in Fig. 2.2.




Figure 2.2: Neutron-induced cross sections of ^{238}U [23].

2.3.2 Phase space

The neutron state is described in the six-dimensional phase space spanned by position, direction, and energy [22]. The position vector is [22]

$$\vec{r} = x\hat{e}_x + y\hat{e}_y + z\hat{e}_z, \quad (2.4)$$

whereas the direction of flight is represented by the unit vector [22]

$$\vec{\Omega} = \frac{\vec{v}}{|\vec{v}|}. \quad (2.5)$$

Its Cartesian components satisfy

$$\vec{\Omega} = \Omega_x\hat{e}_x + \Omega_y\hat{e}_y + \Omega_z\hat{e}_z, \quad \Omega_x^2 + \Omega_y^2 + \Omega_z^2 = 1. \quad (2.6)$$

Introducing the polar angle θ , the azimuthal angle ϕ , and the directional cosine $\mu = \cos \theta$, one obtains

$$\begin{aligned} \Omega_z &= \cos \theta = \mu, \\ \Omega_x &= \sin \theta \cos \phi = \sqrt{1 - \mu^2} \cos \phi, \\ \Omega_y &= \sin \theta \sin \phi = \sqrt{1 - \mu^2} \sin \phi. \end{aligned} \quad (2.7)$$




Figure 2.3: Position and direction of flight of a neutron [15].

The solid angle element is therefore

$$d\Omega = \sin \theta \, d\theta \, d\phi = d\phi \, d\mu, \quad (2.8)$$

with $\phi \in [0, 2\pi)$ and $\mu \in [-1, 1]$ [22, 21].

The full phase-space differential element is $d^3r \, d\Omega \, dE$. This phase space is the natural domain of the neutron distribution and of the transport equation. For time-dependent problems, the neutron distribution evolves in time; the time variable will be denoted by t .

2.3.3 Neutron cross sections

A cross section quantifies the likelihood that a neutron undergoes a specified interaction with a target nucleus [16, 23]. The *microscopic cross section* $\sigma_{x,i}(E)$ refers to reaction type x with isotope i and is measured in barns, with 1 barn = 10^{-24} cm²; it depends on the incident neutron energy. For a material composed of several isotopes, the corresponding *macroscopic cross section* is obtained by summing over the isotopic number densities N_i ,

$$\Sigma_x(\vec{r}, E, t) = \sum_i N_i(\vec{r}, t) \sigma_{x,i}(E), \quad (2.9)$$

and has units of cm⁻¹ [22]. It can be interpreted as the interaction probability per unit path length.

The *total cross section* is the sum of all possible reaction channels,

$$\Sigma_t = \Sigma_s + \Sigma_a = \Sigma_s + \Sigma_f + \Sigma_c + \cdots, \quad (2.10)$$

where Σ_s is the scattering cross section, Σ_a the absorption cross section, Σ_f the fission cross section and Σ_c the capture cross section.

The scattering process is described more precisely by the double differential scattering kernel

$$\Sigma_s(\vec{r}, \vec{\Omega}' \rightarrow \vec{\Omega}, E' \rightarrow E, t) \, d\Omega \, dE, \quad (2.11)$$

which gives the probability per unit path length that a neutron at $(\vec{r}, \vec{\Omega}', E')$ scatters into $(\vec{r}, \vec{\Omega}, E)$ about $d\Omega \, dE$ [22]. In an isotropic homogeneous medium, the angular dependence enters only through the cosine of the scattering angle,

$$\mu_0 \equiv \cos \theta_0 = \vec{\Omega} \cdot \vec{\Omega}', \quad (2.12)$$

so that one may equivalently write [22]

$$\Sigma_s(\vec{r}, \vec{\Omega}' \rightarrow \vec{\Omega}, E' \rightarrow E, t) \equiv \Sigma_s(\vec{r}, \mu_0, E' \rightarrow E, t). \quad (2.13)$$

The fission and capture cross sections, $\Sigma_f(\vec{r}, E, t)$ and $\Sigma_c(\vec{r}, E, t)$, quantify the probabilities per unit path length that a neutron induces fission or is absorbed without fission, respectively [16].

2.3.4 Fission source parameters

This section introduces the quantities that characterise the neutron source generated by fission and that will appear repeatedly in the transport and noise formulations.

The *prompt fission spectrum* $\chi_p(E)$ is the probability density function of the emission energy of prompt fission neutrons and satisfies [16, 22]

$$\int_0^\infty \chi_p(E) dE = 1. \quad (2.14)$$

Similarly, delayed neutrons emitted by precursor family q are characterised by a delayed fission spectrum $\chi_{d,q}(E)$ [22].

The average number of neutrons emitted per fission, $\nu(E)$, depends on the energy of the neutron that induces fission. It is decomposed as [20]

$$\nu(E) = \nu_p(E) + \nu_d(E), \quad (2.15)$$

where $\nu_p(E)$ and $\nu_d(E)$ are the prompt and delayed contributions. The delayed neutron fraction is then [20]

$$\beta(E) = \frac{\nu_d(E)}{\nu(E)}. \quad (2.16)$$

Delayed neutrons are commonly represented by grouping the precursors into Q families indexed by q . Each family is characterised by a delayed fraction $\beta_q(E)$, a decay constant λ_q , and a delayed spectrum $\chi_{d,q}(E)$, with [22]

$$\sum_{q=1}^Q \beta_q(E) = \beta(E). \quad (2.17)$$

The corresponding precursor concentration is denoted by $C_q(\vec{r}, t)$ and has units of cm^{-3} [22].

Equivalently, the delayed contribution of family q may be written as

$$\nu_{d,q}(E) = \beta_q(E) \nu(E), \quad (2.18)$$

so that

$$\nu_d(E) = \sum_{q=1}^Q \nu_{d,q}(E). \quad (2.19)$$

2.3.5 Neutron flux and current

Let $n(\vec{r}, \vec{\Omega}, E, t)$ be the angular neutron density, defined so that

$$n(\vec{r}, \vec{\Omega}, E, t) d^3r d\Omega dE \quad (2.20)$$

is the expected number of neutrons at time t in $d^3r d\Omega dE$ about $(\vec{r}, \vec{\Omega}, E)$. Let $v(E)$ be the neutron speed. The *angular flux* is then defined as [15]

$$\varphi(\vec{r}, \vec{\Omega}, E, t) = n(\vec{r}, \vec{\Omega}, E, t) v(E), \quad (2.21)$$

with units $\text{cm}^{-2} \text{s}^{-1} \text{sr}^{-1} \text{eV}^{-1}$. Physically, φ measures the neutron track length travelled per unit volume, per unit solid angle, per unit energy and per unit time [15].

The corresponding *scalar flux* is obtained by integrating over all directions [15],

$$\Phi(\vec{r}, E, t) = \int_{4\pi} \varphi(\vec{r}, \vec{\Omega}, E, t) d\Omega, \quad (2.22)$$

while the *angular current* is

$$\vec{j}(\vec{r}, \vec{\Omega}, E, t) = \vec{\Omega} \varphi(\vec{r}, \vec{\Omega}, E, t), \quad (2.23)$$

and the *scalar current* is

$$\vec{J}(\vec{r}, E, t) = \int_{4\pi} \vec{j}(\vec{r}, \vec{\Omega}, E, t) d\Omega. \quad (2.24)$$

Throughout this thesis, the notation is kept distinct as follows: φ denotes angular flux, Φ denotes scalar flux in transport theory, and ϕ will later denote the one-group scalar flux in diffusion theory.

2.4 Neutron transport equation

The time-dependent neutron transport equation, or linear Boltzmann equation, expresses the local balance between neutron production, streaming, removal and source terms [22]. For a multiplying system with Q delayed neutron precursor families and no external source, it reads

$$\begin{aligned} & \left(\frac{1}{v(E)} \frac{\partial}{\partial t} + \vec{\Omega} \cdot \vec{\nabla} + \Sigma_t(\vec{r}, E, t) \right) \varphi(\vec{r}, \vec{\Omega}, E, t) = \\ & \int_0^{+\infty} \int_{4\pi} \Sigma_s(\vec{r}, \vec{\Omega}' \rightarrow \vec{\Omega}, E' \rightarrow E, t) \varphi(\vec{r}, \vec{\Omega}', E', t) d\Omega' dE' \\ & + \frac{\chi_p(E)}{4\pi} \int_0^{+\infty} \nu_p(E') \Sigma_f(\vec{r}, E', t) \Phi(\vec{r}, E', t) dE' \\ & + \frac{1}{4\pi} \sum_{q=1}^Q \chi_{d,q}(E) \lambda_q C_q(\vec{r}, t). \end{aligned} \quad (2.25)$$

The precursor concentrations satisfy

$$\frac{\partial C_q(\vec{r}, t)}{\partial t} = \int_0^{+\infty} \nu_{d,q}(E') \Sigma_f(\vec{r}, E', t) \Phi(\vec{r}, E', t) dE' - \lambda_q C_q(\vec{r}, t). \quad (2.26)$$

The terms in Eq. (2.25) have a direct physical interpretation [22]:

- i) the time derivative term describes the temporal accumulation of neutrons;
- ii) the streaming term accounts for neutron motion through space;
- iii) the collision term represents removal from the current phase space state;
- iv) the scattering source term describes neutrons scattered into $(\vec{r}, \vec{\Omega}, E)$ from other directions and energies;
- v) the prompt fission source term represents prompt neutrons born at energy E ;
- vi) the delayed neutron source term accounts for emission from precursor decay.

Equation (2.26) contains the corresponding balance for each precursor family: production by fission and loss by radioactive decay.

In the following sections, explicit dependencies on $(\vec{r}, \vec{\Omega}, E, t)$ will be omitted whenever no ambiguity arises.

2.4.1 Operator form of the transport equation

The precursor equation can be solved formally. For an initial time t_0 , one has [21, 22]

$$C_q(\vec{r}, t) = e^{-\lambda_q(t-t_0)} C_q(\vec{r}, t_0) + \int_{t_0}^t e^{-\lambda_q(t-t')} \int_0^{+\infty} \nu_{d,q}(E') \Sigma_f(\vec{r}, E', t') \Phi(\vec{r}, E', t') dE' dt'. \quad (2.27)$$

If t_0 is taken sufficiently far in the past, the exponentially decaying initial term may be neglected. Inserting the remaining contribution into the transport equation leads to a concise operator form, namely,

$$\left(\frac{1}{v} \frac{\partial}{\partial t} + L - S - F \right) \varphi = 0, \quad (2.28)$$

where

$$L\varphi = \vec{\Omega} \cdot \vec{\nabla} \varphi + \Sigma_t \varphi, \quad (2.29)$$

$$S\varphi = \int_0^{+\infty} \int_{4\pi} \Sigma_s \varphi d\Omega' dE', \quad (2.30)$$

$$F\varphi = F_p\varphi + F_d\varphi. \quad (2.31)$$

The prompt and delayed fission source operators are defined as

$$F_p\varphi = \frac{\chi_p}{4\pi} \int_0^{+\infty} \nu_p \Sigma_f \Phi dE', \quad (2.32)$$

$$F_d\varphi = \frac{1}{4\pi} \sum_{q=1}^Q \chi_{d,q} \lambda_q \int_{t_0}^t e^{-\lambda_q(t-t')} \int_0^{+\infty} \nu_{d,q} \Sigma_f \Phi dE' dt'. \quad (2.33)$$

In this notation, the time dependent Boltzmann operator is

$$\mathcal{B} = \frac{1}{v} \frac{\partial}{\partial t} + L - S - F, \quad \mathcal{B}\varphi = 0. \quad (2.34)$$

This operator notation is particularly convenient in neutron noise theory, where perturbations are introduced as modifications of the underlying transport operator [12, 13].

2.4.2 Steady-state transport equation

In a multiplying medium, a non-trivial time-independent solution does not necessarily exist for the physical system as it stands. The effective multiplication factor k_{eff} is therefore introduced to rescale the fission production term and ensure the existence of a stationary reference solution. The steady state transport equation is then written as the k eigenvalue problem

$$(L - S) \varphi = \frac{1}{k_{\text{eff}}} F \varphi, \quad (2.35)$$

where k_{eff} is the effective multiplication factor introduced in Section 2.2.1. Solving this equation yields the fundamental neutron flux shape and the corresponding criticality condition.

In the remainder of the thesis, the steady state solution will often be referred to as the *critical flux* and denoted by φ_c .

2.5 Multi-group approximation

In the continuous-energy formulation, neutron energy is treated as a continuous variable. In deterministic reactor calculations, the energy variable is discretised

into a finite number of intervals, called *energy groups*, indexed by $g = 0, 1, \dots, G$ [15]. Group g corresponds to the interval $[E_g, E_{g-1}]$, where E_0 is the highest energy boundary and E_G the lowest. This discretisation is sketched in Fig. 2.4.




Figure 2.4: Energy group discretisation [15].

The multi-group approximation replaces the continuous-energy dependence by a set of coupled equations, one for each group. This retains the most important spectral information, while making the problem tractable for deterministic solvers. Note that for Monte Carlo solvers the energy discretization is not necessary; nonetheless, for special applications and during code-to-code comparisons, it might be useful to solve by Monte Carlo simulation multi-group problems, in order to reduce the memory burden and the computing cost. In the following, we will adopt this strategy for the benchmark comparisons of the newly developed neutron noise solvers.

2.5.1 Useful definitions

For compactness, the integral over energy group g is denoted by

$$\int_g dE \equiv \int_{E_g}^{E_{g-1}} dE. \quad (2.36)$$

The *group angular flux* and *group scalar flux* are defined as [15]

$$\varphi^g(\vec{r}, \vec{\Omega}, t) = \int_g \varphi(\vec{r}, \vec{\Omega}, E, t) dE, \quad (2.37)$$

$$\Phi^g(\vec{r}, t) = \int_g \Phi(\vec{r}, E, t) dE. \quad (2.38)$$

The group-integrated prompt and delayed fission spectra are [15]

$$\chi_p^g = \int_g \chi_p(E) dE, \quad \chi_{d,q}^g = \int_g \chi_{d,q}(E) dE. \quad (2.39)$$

Rather than defining ν^g alone, it is more useful in the transport equation to introduce group-averaged reaction coefficients such as [15]

$$(\nu \Sigma_f)^g(\vec{r}, t) = \frac{\int_g \nu(E) \Sigma_f(\vec{r}, E, t) \Phi(\vec{r}, E, t) dE}{\int_g \Phi(\vec{r}, E, t) dE}, \quad (2.40)$$

and similarly, for any reaction type x [15],

$$\Sigma_x^g(\vec{r}, t) = \frac{\int_g \Sigma_x(\vec{r}, E, t) \Phi(\vec{r}, E, t) dE}{\int_g \Phi(\vec{r}, E, t) dE}. \quad (2.41)$$

In practice, the quantity $\Phi(\vec{r}, E, t)$ is not known in advance, since it is precisely the target of the calculations. Some guess function is therefore introduced instead of the true solution. The weighted quantities in Eqs. (2.41) or (2.40) are thus effective group constants (i.e. they are not basic nuclear data) and depend on the weighting spectrum chosen to collapse the continuous-energy dependence.

The *group-to-group scattering cross section* is defined as [15]

$$\Sigma_s^{g' \rightarrow g}(\vec{r}, \vec{\Omega}' \rightarrow \vec{\Omega}, t) = \frac{\int_g \int_{g'} \Sigma_s(\vec{r}, \vec{\Omega}' \rightarrow \vec{\Omega}, E' \rightarrow E, t) \varphi(\vec{r}, \vec{\Omega}', E', t) dE' dE}{\int_{g'} \varphi(\vec{r}, \vec{\Omega}', E', t) dE'}. \quad (2.42)$$

The group-averaged delayed fraction of precursor family q is defined consistently as

$$\beta_q^g(\vec{r}, t) = \frac{\int_g \beta_q(E) \nu(E) \Sigma_f(\vec{r}, E, t) \Phi(\vec{r}, E, t) dE}{\int_g \nu(E) \Sigma_f(\vec{r}, E, t) \Phi(\vec{r}, E, t) dE}. \quad (2.43)$$

Finally, the inverse speed term is replaced by a group-averaged quantity [15],

$$\frac{1}{v^g} = \frac{\int_g \frac{1}{v(E)} \Phi(\vec{r}, E, t) dE}{\int_g \Phi(\vec{r}, E, t) dE}. \quad (2.44)$$

2.5.2 Multi-group transport equation

Integrating Eq. (2.25) over group g and replacing the continuous-energy integrals by sums over groups yields the time-dependent multi-group transport equation [22]

$$\begin{aligned} & \left(\frac{1}{v^g} \frac{\partial}{\partial t} + \vec{\Omega} \cdot \vec{\nabla} + \Sigma_t^g(\vec{r}, t) \right) \varphi^g(\vec{r}, \vec{\Omega}, t) = \\ & \sum_{g'=0}^G \int_{4\pi} \Sigma_s^{g' \rightarrow g}(\vec{r}, \vec{\Omega}' \rightarrow \vec{\Omega}, t) \varphi^{g'}(\vec{r}, \vec{\Omega}', t) d\Omega' \\ & + \frac{\chi_p^g}{4\pi} \sum_{g'=0}^G \left(1 - \sum_{q=1}^Q \beta_q^{g'} \right) (\nu \Sigma_f)^{g'}(\vec{r}, t) \Phi^{g'}(\vec{r}, t) \\ & + \frac{1}{4\pi} \sum_{q=1}^Q \chi_{d,q}^g \lambda_q C_q(\vec{r}, t), \end{aligned} \quad (2.45)$$

with precursor equations

$$\frac{\partial C_q(\vec{r}, t)}{\partial t} = \sum_{g'=0}^G \beta_q^{g'} (\nu \Sigma_f)^{g'}(\vec{r}, t) \Phi^{g'}(\vec{r}, t) - \lambda_q C_q(\vec{r}, t). \quad (2.46)$$

Equations (2.45) and (2.46) form a coupled system of G transport equations and Q precursor equations. The coupling arises from scattering, from fission neutrons born in one group and emitted into another, and from delayed neutron production.

2.6 Kinetic approximations

Starting from the time-dependent neutron transport problem, reactor kinetic approximations are commonly obtained by factorising the neutron flux into an amplitude function and a shape function. In the following, this idea is written in a one-group scalar-flux notation for simplicity, since this is sufficient for the interpretation of the frequency regimes used later in the thesis. The same factorisation procedure can be formulated directly in transport theory by using the angular flux and by introducing effective quantities through appropriate direct-adjoint flux weighting [1, 20].

The flux is written as

$$\phi(\vec{r}, t) = P(t) \psi(\vec{r}, t), \quad (2.47)$$

where $P(t)$ represents the neutron population amplitude and $\psi(\vec{r}, t)$ its spatial shape. This decomposition is not unique unless a normalisation condition is imposed. In the general transport setting, this condition involves the adjoint flux of a reference critical system. In the one-group diffusion setting considered here for simplicity, which is self-adjoint under standard assumptions, the condition may be written as [1]

$$\frac{\partial}{\partial t} \int_V \phi_0(\vec{r}) \psi(\vec{r}, t) d\vec{r} = 0, \quad (2.48)$$

where $\phi_0(\vec{r})$ denotes the static flux of the reference critical configuration,

$$\phi_0(\vec{r}) = P_0 \psi(\vec{r}, 0). \quad (2.49)$$

Substituting Eq. (2.47) into the one-group kinetic equations and applying the normalisation condition leads to the point-kinetics equations [1, 20]

$$\begin{cases} \frac{dP(t)}{dt} = \frac{\rho(t) - \beta}{\Lambda(t)} P(t) + \lambda C(t), \\ \frac{dC(t)}{dt} = \frac{\beta}{\Lambda(t)} P(t) - \lambda C(t). \end{cases} \quad (2.50)$$

where $\rho(t)$ is the reactivity, β is the delayed neutron fraction, $\Lambda(t)$ is the effective neutron generation time, $C(t)$ is the precursor amplitude, and λ is the precursor decay constant. In a more general multi-group or transport formulation, ρ , β , and Λ are replaced by effective quantities obtained through suitable weighting with the direct and adjoint fluxes [20, 1].

At the same time, the shape function satisfies an equation of the form [1]

$$\frac{1}{v} \left[\frac{\partial \psi(\vec{r}, t)}{\partial t} + \frac{1}{P(t)} \frac{dP(t)}{dt} \psi(\vec{r}, t) \right] = D \nabla^2 \psi(\vec{r}, t) + [\nu \Sigma_f(\vec{r}, t)(1 - \beta) - \Sigma_a(\vec{r}, t)] \psi(\vec{r}, t) + \frac{Q_d(\vec{r}, t)}{P(t)}, \quad (2.51)$$

where $Q_d(\vec{r}, t)$ is the delayed neutron source term. The difficulty of solving this equation motivates additional approximations on the shape function.

In the *point reactor approximation*, the shape is assumed to remain equal to the static one,

$$\psi(\vec{r}, t) = \phi_0(\vec{r}), \quad \forall t, \quad (2.52)$$

so that

$$\phi(\vec{r}, t) = P(t) \phi_0(\vec{r}). \quad (2.53)$$

This approximation is appropriate when the perturbation causes almost no spatial distortion of the flux and the time dependence can be represented mainly by a global amplitude variation [20, 1].

In the *adiabatic approximation*, the shape is allowed to follow the instantaneous material state, but it is assumed to do so sufficiently slowly that, at each time, it satisfies a static eigenvalue problem,

$$D \nabla^2 \psi(\vec{r}, t) + \left[\frac{\nu \Sigma_f(\vec{r}, t)}{k(t)} - \Sigma_a(\vec{r}, t) \right] \psi(\vec{r}, t) = 0, \quad (2.54)$$

together with the normalisation condition of Eq. (2.48) [1]. This approximation is more flexible than the point reactor approximation, because it allows the flux shape to evolve with the instantaneous configuration, but it remains limited to sufficiently slow and moderate perturbations.

These kinetic approximations are historically important because they separate the global amplitude response from the spatial redistribution of the neutron flux. They also provide the basis for the interpretation of the zero-power reactor transfer function and of the frequency regimes discussed later in the thesis. In the present work, however, they are used only as interpretive tools: the neutron noise calculations themselves are performed with the transport-based linear and non-perturbative formulations introduced in Chapter 4.

Chapter 3

Monte Carlo Method

The Monte Carlo method is a class of numerical techniques in which the solution of a mathematical or physical problem is obtained through the statistical analysis of random samples [24, 25]. Developed during the early years of nuclear research, it has become one of the main computational tools of reactor physics because it can represent neutron transport in a very direct way, with relatively few physical approximations [24, 22].

In neutron transport applications, the Monte Carlo method simulates the stochastic evolution of individual particle histories. Each neutron undergoes a sequence of free flights and interactions sampled from the underlying probability laws, and macroscopic observables such as fluxes, reaction rates and multiplication factors are estimated from the accumulated scores of many such histories [22, 24]. In this sense, the method provides a statistical solution of the Boltzmann equation, in contrast with deterministic approaches based on discretisation in space, angle and energy.

One advantage of Monte Carlo methods is their flexibility. Complex geometries, detailed cross section data and intricate interaction physics can be incorporated with limited additional modelling effort. The main drawback is statistical uncertainty: the accuracy of the estimated quantities improves only slowly with the number of simulated histories, so high precision may require substantial computational cost [22]. This issue is particularly relevant in neutron transport problems involving strong attenuation, localised tallies or weakly contributing events.

The present chapter recalls the statistical concepts that underpin Monte Carlo methods and introduces the basic ingredients of Monte Carlo neutron transport. Particular attention is paid to particle histories, estimators and eigenvalue calculations, since these notions are central to the frequency-domain noise formulations

discussed in the following chapters.

3.1 Statistical foundations

Monte Carlo methods rely on the interpretation of physical observables as expectations of random quantities. Only a limited set of statistical concepts is needed in the present context: random variables, expectation, variance, and the sampling properties of estimators.

3.1.1 Random variables, distributions and moments

A *random variable* ξ is a mathematical function capable of representing the outcome of a random experiment [26]. In Monte Carlo transport, examples include the distance to the next collision, the number of collisions in a region, or the score associated with a particle history.

The probabilistic behaviour of a real-valued random variable ξ is described by its cumulative distribution function [25, 26]

$$F_\xi(t) = P(\xi \leq t), \quad (3.1)$$

which gives the probability that ξ takes a value not larger than t . If F_ξ is differentiable, its derivative [26]

$$f_\xi(t) = \frac{dF_\xi}{dt} \quad (3.2)$$

is the corresponding probability density function. For discrete variables, the distribution is instead described by point probabilities $P(\xi = x_i)$.

The first quantity of interest is the *expected value*, or mean, denoted by $\mathbb{E}[\xi]$. For a discrete random variable [25, 26],

$$\mathbb{E}[\xi] = \sum_i x_i P(\xi = x_i), \quad (3.3)$$

whereas for a continuous random variable with a probability density function [25]

$$\mathbb{E}[\xi] = \int_{-\infty}^{+\infty} t f_\xi(t) dt. \quad (3.4)$$

More generally, the n th moment is [26]

$$\mu_n = \mathbb{E}[\xi^n]. \quad (3.5)$$

A second essential quantity is the *variance* [25],

$$\text{Var}[\xi] = \mathbb{E}[(\xi - \mathbb{E}[\xi])^2], \quad (3.6)$$

which measures the spread of the distribution around its mean. It is often convenient to use the equivalent identity [25, 26]

$$\text{Var}[\xi] = \mathbb{E}[\xi^2] - (\mathbb{E}[\xi])^2. \quad (3.7)$$

The square root of the variance,

$$\sigma = \sqrt{\text{Var}[\xi]}, \quad (3.8)$$

is the standard deviation and has the same physical units as the random variable itself [25].

In Monte Carlo calculations, the quantity of interest is often an expectation $\mathbb{E}[\xi]$ that cannot be evaluated analytically. The role of the simulation is then to construct an estimator of this expectation from a finite sample of histories.

3.1.2 Sample mean and statistical uncertainty

Let $\xi_1, \xi_2, \dots, \xi_N$ be N realisations of the random variable ξ , generated under the same sampling law, i.e. assuming that the ξ_i are independent and identically distributed, or IID. The basic Monte Carlo estimator of the mean is the sample average [25]

$$\bar{\xi}^{(N)} = \frac{1}{N} \sum_{i=1}^N \xi_i. \quad (3.9)$$

By construction, $\mathbb{E}[\bar{\xi}^{(N)}] = \mathbb{E}[\xi]$ [24]. Its variance is [25]

$$\text{Var} [\bar{\xi}^{(N)}] = \frac{\sigma^2}{N}, \quad (3.10)$$

so that the standard deviation of the estimator decreases as

$$\sigma_{\bar{\xi}^{(N)}} = \frac{\sigma}{\sqrt{N}}. \quad (3.11)$$

This $N^{-1/2}$ convergence is one of the defining features of Monte Carlo methods. In practical terms, reducing the standard deviation by a factor of two requires roughly four times more histories.

A fundamental result underlying probability theory is the law of large numbers, which ensures that the sample mean converges almost surely to the true expectation as N increases. In addition, the central limit theorem states that, under broad conditions, the distribution of the centred and normalised sample mean approaches a Gaussian law for large N . This justifies the common use of standard errors and confidence intervals in Monte Carlo analysis [26].

In transport calculations, the independence assumption is exact only for genuinely independent histories. In fixed-source simulations this is usually a correct assumption, whereas in eigenvalue calculations correlations are introduced by the fact that the source of a given generation corresponds to the fission events of the previous generation. This point will be revisited later in the chapter.

3.1.3 Sample variance

Since the true variance σ^2 is generally unknown, it must also be estimated from the simulated sample. The standard unbiased estimator is the sample variance

$$s_N^2 = \frac{1}{N-1} \sum_{i=1}^N (\xi_i - \bar{\xi}^{(N)})^2, \quad (3.12)$$

which satisfies

$$\mathbb{E}[s_N^2] = \sigma^2 \quad (3.13)$$

under the assumption of IID samples [24]. The factor $(N-1)^{-1}$ is the classical Bessel correction [27].

From Eq. (3.12), the estimated variance of the sample mean is [24]

$$\widehat{\text{Var}} [\bar{\xi}^{(N)}] = \frac{s_N^2}{N}, \quad (3.14)$$

and the corresponding estimated standard error is

$$\hat{\sigma}_{\bar{\xi}^{(N)}} = \frac{s_N}{\sqrt{N}}. \quad (3.15)$$

These quantities are the basic statistical outputs reported in Monte Carlo calculations. They measure the uncertainty due to finite sampling and provide the basis for confidence intervals and for the statistical tests used later in the thesis when comparing different estimators and methods.

3.1.4 Bias and figure of merit

An estimator $\hat{X}$ of $\mathbb{E}[\xi]$ is *unbiased* if $\mathbb{E}[\hat{X}] = \mathbb{E}[\xi]$. However, unbiasedness alone is not sufficient to judge its practical quality. Two unbiased estimators may have very different variances, and therefore very different efficiencies.

A common performance metric is the *figure of merit*, conventionally defined as [24]

$$\text{FoM} = \frac{1}{\sigma_X^2 t_c}, \quad (3.16)$$

where σ_X^2 is the variance of the estimator and t_c is the computational time. A larger figure of merit corresponds to a more efficient estimator. In practical Monte Carlo work, improving efficiency means reducing variance without increasing cost too much, or reducing cost while maintaining the same uncertainty.

3.2 Monte Carlo solution of the transport equation

3.2.1 Particle histories and random walks

In neutron transport, the Monte Carlo method replaces the direct solution of the Boltzmann equation by the simulation of a large number of particle histories. A *history* is the full random walk of one particle and all its descendants, from birth to termination. Each history is a succession of free flights and collisions, determined by random sampling from the physical probability laws [26].

Macroscopic observables are not obtained from a single history, but from averages over many histories. In this way, transport quantities are interpreted as expectations of suitable random scores.

The random sampling itself is driven by pseudo random numbers uniformly distributed in $[0,1)$, which are transformed into samples from the physical distributions of interest by techniques such as inverse transform sampling, composition methods or rejection sampling [25, 26].

3.2.2 Sampling of free flight and collision type

A key step of the simulation is the sampling of the distance to the next collision. In a homogeneous medium with total macroscopic cross section Σ_t , the probability that a neutron travels a distance s greater than ℓ without collision is [20, 24]

$$P(s > \ell) = e^{-\Sigma_t \ell}. \quad (3.17)$$

Hence, the free flight distance is exponentially distributed and can be sampled as [24, 25]

$$s = -\frac{\ln \eta}{\Sigma_t}, \quad \eta \sim \mathcal{U}[0,1). \quad (3.18)$$

In heterogeneous media or in the presence of spatially-varying cross sections, this simple form is replaced by more elaborate tracking or rejection techniques, but the underlying principle remains the same.

Once a collision occurs, the reaction type is sampled from the partial cross sections [20, 28]. If the possible outcomes are scattering, capture and fission, their

probabilities are

$$P_s = \frac{\Sigma_s}{\Sigma_t}, \quad P_c = \frac{\Sigma_c}{\Sigma_t}, \quad P_f = \frac{\Sigma_f}{\Sigma_t}. \quad (3.19)$$

After the reaction type has been selected, the outgoing particle state is sampled from the corresponding physical law. For scattering, this means sampling the outgoing energy and direction from the scattering kernel. For fission, it means generating the emitted neutrons according to the multiplicity and fission spectrum.

3.2.3 Analog and non-analog transport

In *analog* Monte Carlo, every event is sampled according to the true physical probabilities, so the simulated process reproduces the physical random walk directly. This has the advantage of conceptual simplicity, but it may be inefficient when the quantity of interest depends on rare events or on trajectories with very low probability [24, 25].

For this reason, practical transport calculations often use *non-analog* techniques, in which the sampling law is modified to improve efficiency and the induced bias is compensated by adjusting particle weights [24, 25]. This is the basis of importance sampling and of classical variance reduction techniques such as splitting, Russian roulette and weight windows.

A common feature of non-analog Monte Carlo is to assign a statistical weight w to each particle, so that the contribution of each particle to the tallies is proportional to its weight. Arbitrary sampling laws can then be used instead of the physical ones, in order to sample relevant events more frequently, provided that the particle weights are modified accordingly. If a particle with a weight w should be moved to a generic phase-space coordinate x with a physical probability density $f(x)$ but an arbitrary probability density $g(x)$ is used instead, then the particle weight must be adjusted as

$$w' = \frac{f(x)}{g(x)} w. \quad (3.20)$$

This rule ensures that the estimators remain unbiased [24]. This weighted formulation is especially important in the later chapters of the thesis, where complex-valued source terms, operators and weights are introduced in frequency-domain neutron noise calculations.

3.3 Estimators in neutron transport

3.3.1 General notion of a tally

A *tally* is the Monte Carlo estimate of a transport observable. Formally, if X denotes the quantity of interest, a Monte Carlo calculation constructs a score ξ_i for each history such that [24]

$$X = \mathbb{E}[\xi], \quad (3.21)$$

and estimates X through

$$\hat{X}_N = \frac{1}{N} \sum_{i=1}^N \xi_i. \quad (3.22)$$

Different estimators correspond to different definitions of the score variable. Their quality is assessed in terms of unbiasedness, variance and computational cost [24].

For neutron transport, the most common quantities of interest are fluxes, reaction rates and multiplication factors. The corresponding tallies can often be constructed in more than one way. This is particularly relevant in the present thesis, because different estimator formulations will later be compared in the context of frequency-domain noise calculations.

3.3.2 Track-length estimator

The track-length estimator is one of the most widely used estimators for the neutron flux [24, 25]. Consider a spatial region of volume V . If a particle travels a track segment of length ℓ inside that region with weight w , its contribution to the scalar flux is proportional to $w\ell$. Summing over all segments of all histories yields the estimator [20]

$$\hat{\Phi}_{TL,V} = \frac{1}{VN} \sum_{i=1}^N \sum_k w_{ik} \ell_{ik}, \quad (3.23)$$

where the inner sum runs over all track segments k of history i lying inside V .

The intuitive interpretation is straightforward: the scalar flux is proportional to the total path length travelled by neutrons per unit volume. In practice, the track-length estimator is often preferred because it uses information continuously along the particle path rather than only at collision sites, and it therefore tends to have lower variance for flux-like quantities. However, it requires estimation of the lengths spent by each particle in each tally region, $\ell_{i,k}$, which may increase the computational cost.

Reaction rates can be estimated in a corresponding way by weighting the track length with the relevant macroscopic cross section. If the material in the spatial

region is homogeneous, for reaction type x ,

$$\hat{R}_{x,TL,V} = \frac{1}{VN} \sum_{i=1}^N \sum_k w_{ik} \Sigma_x(\vec{r}_{ik}, E_{ik}) \ell_{ik}. \quad (3.24)$$

3.3.3 Collision estimator

An alternative flux estimator uses collision events rather than track segments [24, 25]. In its standard form, each collision event c of history i occurring inside region V contributes a score proportional to w/Σ_t , so that

$$\hat{\Phi}_{C,V} = \frac{1}{VN} \sum_{i=1}^N \sum_c \frac{w_{ic}}{\Sigma_t(\vec{r}_{ic}, E_{ic})}. \quad (3.25)$$

This estimator is also unbiased. The corresponding collision estimator for the reaction rate of type x is

$$\hat{R}_{x,C,V} = \frac{1}{VN} \sum_{i=1}^N \sum_c w_{ic} \frac{\Sigma_x(\vec{r}_{ic}, E_{ic})}{\Sigma_t(\vec{r}_{ic}, E_{ic})}. \quad (3.26)$$

Compared with the track-length estimator, the collision estimator uses only information at collision points. Depending on the problem, one or the other may be more efficient. This distinction becomes especially important in the later chapters, where the treatment of the frequency-dependent term in the noise equation leads to different estimator formulations and where the consistency of the track-length and collision approaches must be examined carefully.

3.4 Monte Carlo eigenvalue calculations

3.4.1 The k -eigenvalue problem

In multiplying systems, the neutron source is not prescribed a priori because it depends on the unknown fission distribution. Monte Carlo criticality calculations therefore solve the steady state transport equation in its eigenvalue form by an iterative procedure often interpreted as a stochastic power method [24].

Starting from an initial guess for the source distribution, one simulates a generation of source neutrons and records the fission neutrons produced by their histories. These fission neutrons define the source for the next generation. Repeating this procedure over many generations drives the source distribution towards the fundamental eigenmode, while the effective multiplication factor is estimated from the ratio between produced and starting source populations [28, 24].

If $W_s^{(m)}$ denotes the total starting source weight in generation m and $W_f^{(m)}$ the total weight of produced fission neutrons, a generation-wise estimator of k_{eff} is

$$k^{(m)} = \frac{W_f^{(m)}}{W_s^{(m)}}. \quad (3.27)$$

The final estimate is obtained by averaging over sufficiently many active generations once the source has converged.

3.4.2 Inactive generations, source convergence and population control

Because the initial source guess is generally not equal to the fundamental mode, the first generations are affected by source convergence transients. It is therefore standard practice to discard an initial set of *inactive generations*, used only to allow the fission source to converge, and to accumulate tallies only over the subsequent *active generations* [28].

In addition, the number of neutrons produced by fission will rapidly diverge or vanish if the system is not critical; moreover, it fluctuates from one generation to the next. In order to keep the computational effort under control, the source population is renormalised to a prescribed size at each generation [28]. This step may involve resampling, splitting or combing source particles. Although essential in practice, population control may slightly complicate the interpretation of the statistical errors.

For this reason, the standard formulas for sample variance should be interpreted with care in eigenvalue calculations. They remain extremely useful in practice, but the underlying histories are not independent. This is one of the reasons why source convergence diagnostics and careful settings for inactive and active generations are important in criticality simulations.

Chapter 4

Neutron Noise Theory

The objective of neutron noise theory is to describe the fluctuations of the neutron field around its stationary state and, in particular, their behaviour in time and frequency. In practice, neutron noise is studied in the frequency domain, where the response of the system can be related to oscillatory perturbations of material properties, geometry, or boundary conditions. This chapter introduces the origins of neutron noise and presents the linear and non-perturbative formulations within the framework of continuous-energy transport theory. Particular attention is given to the eigenvalue formulation of the non-perturbative problem, including the reformulated non-perturbative eigenvalue problem later adopted for the one-dimensional vibrating rod benchmark.

4.1 Origins of neutron noise

Neutron noise denotes fluctuations of the neutron population. These fluctuations contain information on the physical state of the system and can therefore be exploited for monitoring, diagnostics, and anomaly detection [1]. In particular, neutron noise measurements can be used to infer properties of the reactor core, such as coolant flow conditions, density variations, or mechanical vibrations of the internals or even of the vessel [1].

Two broad classes of neutron fluctuations are traditionally distinguished according to their physical origin and field of application: *zero-power reactor noise* and *power reactor noise* [1]. The present thesis is concerned exclusively with power reactor noise, but the distinction is briefly recalled here for completeness.

Zero power reactor noise. At low neutron flux, such as in subcritical assemblies, startup conditions, or research reactors operated at very low power, fluctuations are mainly due to the stochastic nature of neutron multiplication itself. Since a single fission event produces several neutrons, correlations are introduced between successive generations, and the counting statistics deviate from a purely Poissonian behaviour. Zero-power noise methods are mainly used for sub-criticality measurements and safeguards applications [1].

Power reactor noise. At power reactor conditions, the neutron population is so large that purely statistical branching fluctuations are comparatively small. The dominant contribution instead comes from physical perturbations of the system, such as vibrations of fuel assemblies or control rods, oscillations of coolant density or temperature, and boiling phenomena. Power reactor noise is therefore primarily interpreted as the neutronic response to external or internal (typically periodic) physical perturbations of the reactor [1].

4.2 Power reactor noise

In power reactor noise theory, time-dependent perturbations are represented as perturbations of the transport operator. In the classical formulation, the perturbation is usually defined with respect to the reference critical steady state of the unperturbed system [1]. In that case, the material properties and the neutron flux are decomposed as

$$\Sigma(\vec{r}, t) = \Sigma_c(\vec{r}) + \delta\Sigma(\vec{r}, t), \quad \varphi(\vec{r}, t) = \varphi_c(\vec{r}) + \delta\varphi(\vec{r}, t), \quad (4.1)$$

where $\Sigma_c(\vec{r})$ and $\varphi_c(\vec{r})$ denote the stationary quantities associated with the reference critical state, while $\delta\Sigma(\vec{r}, t)$ and $\delta\varphi(\vec{r}, t)$ represent the time-dependent perturbation and the induced neutron noise, respectively [1].

For vibration-induced noise, however, this choice is not always the most suitable one. Mechanical vibrations may produce a non-zero time-averaged effect on the material distribution and, consequently, on the reactivity of the system [12]. For this reason, Sanchez proposed a formulation based on a time-averaged, or new steady-state, reference operator, in which the perturbation is defined with respect to the average perturbed system rather than with respect to the initial unperturbed one [12]. This is the formulation adopted in the present work. Accordingly, the material properties and the flux are written as

$$\Sigma(\vec{r}, t) = \Sigma_0(\vec{r}) + \delta\Sigma(\vec{r}, t), \quad \varphi(\vec{r}, t) = \varphi_0(\vec{r}) + \delta\varphi(\vec{r}, t), \quad (4.2)$$

where Σ_0 and φ_0 denote the quantities associated with the time-averaged reference state, while $\delta\Sigma$ and $\delta\varphi$ represent the corresponding fluctuations.

Let $\mathcal{B}_0$ denote the time-independent Boltzmann operator associated with this time-averaged reference state. The reference flux satisfies

$$\mathcal{B}_0 \varphi_0 = 0, \quad (4.3)$$

whereas the perturbed operator is written as

$$\mathcal{B}(t) = \mathcal{B}_0 + \delta\mathcal{B}(t). \quad (4.4)$$

Substituting Eqs. (4.2) into the transport equation and subtracting the reference balance leads to the exact perturbed equation [13]

$$[\mathcal{B}_0 + \delta\mathcal{B}(t)] \delta\varphi(t) = -\delta\mathcal{B}(t) \varphi_0. \quad (4.5)$$

This equation is the starting point for the two main formulations used in this work. The first is the conventional *linear*, or *orthodox*, theory, obtained by neglecting products of perturbations. The second is the novel *non-perturbative* theory, in which the coupling between perturbation and fluctuating flux is fully retained.

4.2.1 Linear theory

The linear theory assumes that the perturbations are small compared with the reference state,

$$|\delta\Sigma| \ll \Sigma_0, \quad (4.6)$$

so that the product $\delta\mathcal{B}(t) \delta\varphi(t)$ is assumed to be of second order and is neglected [1]. Equation (4.5) then reduces to [1]

$$\mathcal{B}_0 \delta\varphi(t) = -\delta\mathcal{B}(t) \varphi_0. \quad (4.7)$$

The noise problem is thus recast as a fixed-source problem, where the unknown is the fluctuating flux $\delta\varphi$ and the source term is generated by the action of the perturbation operator $\delta\mathcal{B}(t)$ on the reference flux [13].

Using the Fourier transform

$$\widehat{f}(\omega) = \int_{-\infty}^{+\infty} e^{-i\omega t} f(t) dt, \quad (4.8)$$

Eq. (4.7) becomes

$$\widehat{\mathcal{B}}_0(\omega) \widehat{\delta\varphi}(\omega) = -\widehat{\delta\mathcal{B}}(\omega) \varphi_0. \quad (4.9)$$

The Fourier transformed reference operator acts on a generic function ψ as

$$\begin{aligned}
 (\widehat{\mathcal{B}}_0(\omega)\psi)(\vec{r}, \vec{\Omega}, E) &= \left(\frac{i\omega}{v(E)} + \vec{\Omega} \cdot \nabla + \Sigma_{t,0}(\vec{r}, E) \right) \psi(\vec{r}, \vec{\Omega}, E) \\
 &- \int_0^{+\infty} \int_{4\pi} \Sigma_{s,0}(\vec{r}, \vec{\Omega}' \rightarrow \vec{\Omega}, E' \rightarrow E) \psi(\vec{r}, \vec{\Omega}', E') d\Omega' dE' \\
 &- \frac{\chi_p(E)}{4\pi} \int_0^{+\infty} \int_{4\pi} \nu_p(E') \Sigma_{f,0}(\vec{r}, E') \psi(\vec{r}, \vec{\Omega}', E') d\Omega' dE' \\
 &- \frac{1}{4\pi} \sum_{q=1}^Q \chi_{d,q}(E) \frac{\lambda_q}{\lambda_q + i\omega} \int_0^{+\infty} \int_{4\pi} \nu_{d,q}(E') \Sigma_{f,0}(\vec{r}, E') \psi(\vec{r}, \vec{\Omega}', E') d\Omega' dE'.
 \end{aligned} \tag{4.10}$$

Similarly, the Fourier transformed perturbation operator is

$$\begin{aligned}
 (\delta\widehat{\mathcal{B}}(\omega)\psi)(\vec{r}, \vec{\Omega}, E) &= \delta\widehat{\Sigma}_t(\vec{r}, E, \omega) \psi(\vec{r}, \vec{\Omega}, E) \\
 &- \int_0^{+\infty} \int_{4\pi} \delta\widehat{\Sigma}_s(\vec{r}, \vec{\Omega}' \rightarrow \vec{\Omega}, E' \rightarrow E, \omega) \psi(\vec{r}, \vec{\Omega}', E') d\Omega' dE' \\
 &- \frac{\chi_p(E)}{4\pi} \int_0^{+\infty} \int_{4\pi} \nu_p(E') \delta\widehat{\Sigma}_f(\vec{r}, E', \omega) \psi(\vec{r}, \vec{\Omega}', E') d\Omega' dE' \\
 &- \frac{1}{4\pi} \sum_{q=1}^Q \chi_{d,q}(E) \frac{\lambda_q}{\lambda_q + i\omega} \int_0^{+\infty} \int_{4\pi} \nu_{d,q}(E') \delta\widehat{\Sigma}_f(\vec{r}, E', \omega) \psi(\vec{r}, \vec{\Omega}', E') d\Omega' dE'.
 \end{aligned} \tag{4.11}$$

Equation (4.9) shows the key property of the orthodox theory: the problem is *decoupled in frequency*, since ω occurs as a fixed parameter in the noise equations. Each Fourier component can be solved independently, because the operator $\widehat{\mathcal{B}}_0(\omega)$ acts only on the unknown at the same frequency ω . The linear theory therefore confines the response to the same frequency content as the imposed perturbation: it cannot represent spectral redistribution among harmonics that may appear in the exact response of a periodically-perturbed reactor.

This limitation is especially relevant for vibration-induced perturbations, for which the exact transport response may exhibit a significant harmonic content beyond the fundamental frequency [12, 13].

4.2.2 Non-perturbative theory

The linear theory neglects the product $\delta\mathcal{B}\delta\varphi$, and is therefore valid only when both the perturbation and the induced response remain sufficiently small. In situations where this product cannot be neglected, the full perturbed transport equation must be retained [13]. Writing the exact flux as $\varphi(t) = \varphi_0 + \delta\varphi(t)$, the governing equation is

$$\mathcal{B}_0\varphi(t) + \delta\mathcal{B}(t)\varphi(t) = 0. \tag{4.12}$$

After Fourier transformation, one obtains [14]

$$\widehat{\mathcal{B}}_0(\omega) \widehat{\varphi}(\omega) + \frac{1}{2\pi} (\widehat{\delta\mathcal{B}} * \widehat{\varphi})(\omega) = 0, \quad (4.13)$$

where

$$(\widehat{\delta\mathcal{B}} * \widehat{\varphi})(\omega) = \int_{-\infty}^{+\infty} \widehat{\delta\mathcal{B}}(\omega - \omega') \widehat{\varphi}(\omega') d\omega'. \quad (4.14)$$

The detailed mathematical properties of the convolution product are recalled in Appendix A.2.

The essential difference with respect to the linear case is that the problem is now globally coupled in frequency. The unknown at frequency ω depends on the solution at all other frequencies through the convolution term. In other words, the non-perturbative (NP) formulation allows a redistribution of spectral content among harmonics and can therefore yield a non-zero response at higher harmonics even when the perturbation is monochromatic.

In explicit form, Eq. (4.13) may be written schematically as [14]

$$\begin{aligned} & \left(\frac{i\omega}{v} + \vec{\Omega} \cdot \nabla + \Sigma_{t,0} \right) \widehat{\varphi}(\omega) - \int_0^{+\infty} \int_{4\pi} \Sigma_{s,0} \widehat{\varphi}(\omega) d\Omega' dE' \\ & - \int_0^{+\infty} \int_{4\pi} \left(\frac{\chi_p \nu_p}{4\pi} + \sum_{q=1}^Q \frac{\chi_{d,q} \nu_{d,q}}{4\pi} \frac{\lambda_q}{\lambda_q + i\omega} \right) \Sigma_{f,0} \widehat{\varphi}(\omega) d\Omega' dE' \\ & + \frac{1}{2\pi} \left[\widehat{\delta\Sigma_t} * \widehat{\varphi} - \int_0^{+\infty} \int_{4\pi} \widehat{\delta\Sigma_s} * \widehat{\varphi} d\Omega' dE' \right] \\ & - \frac{1}{2\pi} \int_0^{+\infty} \int_{4\pi} \left(\frac{\chi_p \nu_p}{4\pi} + \sum_{q=1}^Q \frac{\chi_{d,q} \nu_{d,q}}{4\pi} \frac{\lambda_q}{\lambda_q + i\omega} \right) \widehat{\delta\Sigma_f} * \widehat{\varphi} d\Omega' dE' = 0, \end{aligned} \quad (4.15)$$

where, for brevity, most dependencies on $(\vec{r}, \vec{\Omega}, E)$ have been omitted.

The non-perturbative approach therefore provides a more complete description of neutron noise, especially for vibration-induced perturbations and other situations in which harmonic coupling is important [5, 13, 14]. Its main drawback is that all harmonic components must be solved simultaneously, which makes the problem significantly more expensive than the decoupled linear formulation. Unlike the orthodox linear equation, Eq. (4.15) is homogeneous. In order to solve it by Monte Carlo power iteration, the non-perturbative frequency-domain problem is recast as an eigenvalue problem [14]. Since the operators involved are in general complex-valued and the convolution term couples different harmonic components, the corresponding Monte Carlo treatment requires specific care. These algorithmic aspects are detailed in Chapter 5.

Eigenvalue formulations of the non-perturbative problem

A first, natural eigenvalue formulation is obtained by keeping the perturbations of the total and scattering operators on the left hand side, while placing the stationary and perturbed fission production terms on the right hand side. With the notation introduced above, this gives

$$\left(\frac{i\omega}{v} + \vec{\Omega} \cdot \nabla + \Sigma_{t,0} - S_0\right) \hat{\varphi} + \frac{1}{2\pi} (\widehat{\delta\Sigma_t} - \widehat{\delta S}) * \hat{\varphi} = \frac{1}{k} \left[F_0 \hat{\varphi} + \frac{1}{2\pi} \widehat{\delta F} * \hat{\varphi} \right]. \quad (4.16)$$

This formulation will be referred to in the following as the original *k-scheme*. The eigenvalue *k* should not be confused with the effective multiplication factor of the stationary reactor. It is the eigenvalue associated with this non-perturbative harmonic formulation and acts as the normalization factor in the corresponding power iteration. The original *k-scheme* was used in the first proof-of-concept non-perturbative Monte Carlo calculations and behaves satisfactorily for homogeneous configurations [14].

In the calculations performed in this work, the original *k-scheme* proved sufficient for the homogeneous benchmarks considered in Chapter 6. However, when applied to the one-dimensional vibrating rod benchmark of Chapter 7, it did not converge for realistic vibration amplitudes. The difficulty is related to the fact that the perturbation may locally induce a strong production effect within a generation, so that the balance between production and removal becomes unfavourable before the source renewal step. In practice, converged results could only be obtained by reducing the vibration amplitude to unrealistically small values.

For this reason, an alternative eigenvalue formulation was adopted. It is obtained by moving all perturbation-induced terms to the right hand side, so that the left hand side contains only the time averaged transport, removal and scattering operators. The resulting problem is

$$\left(\frac{i\omega}{v} + \vec{\Omega} \cdot \nabla + \Sigma_{t,0} - S_0\right) \hat{\varphi} = \frac{1}{k'} \left[F_0 \hat{\varphi} + \frac{1}{2\pi} (\widehat{\delta F} + \widehat{\delta S} - \widehat{\delta\Sigma_t}) * \hat{\varphi} \right]. \quad (4.17)$$

This formulation will be referred to as the reformulated *k'-scheme*. As for *k*, the eigenvalue *k'* should not be interpreted as the physical multiplication factor of the reactor. It is instead a numerical eigenvalue associated with the reformulated non-perturbative problem. The prime is used to distinguish it from the eigenvalue *k* of Eq. (4.16), since the two quantities correspond to different eigenvalue formulations of the same homogeneous non-perturbative equation. When the eigenvalue is equal to one, the two formulations are algebraically equivalent, but they lead to different Monte Carlo power-iteration algorithms.

The practical Monte Carlo interpretation of Eqs. (4.16) and (4.17), including the treatment of particles in the current and next generations, is described in Chapter 5.

4.2.3 Point-kinetics interpretation and zero power transfer function

The classical reactor kinetics approximations described in Section 2.6 rely on the factorisation

$$\phi(\vec{r}, t) = P(t) \psi(\vec{r}, t), \quad (4.18)$$

where $P(t)$ is a global amplitude and $\psi(\vec{r}, t)$ is a normalised shape function [1]. In the stationary reference state,

$$\phi_0(\vec{r}) = P_0 \psi_0(\vec{r}). \quad (4.19)$$

Under a small perturbation,

$$P(t) = P_0 + \delta P(t), \quad \psi(\vec{r}, t) = \psi_0(\vec{r}) + \delta \psi(\vec{r}, t), \quad (4.20)$$

and retaining only first order terms gives the classical decomposition of the neutron noise [1]

$$\delta \phi(\vec{r}, t) = \phi_0(\vec{r}) \frac{\delta P(t)}{P_0} + P_0 \delta \psi(\vec{r}, t). \quad (4.21)$$

The first term is a point-kinetics contribution proportional to the stationary shape, whereas the second term represents the spatial distortion of the flux shape caused by the perturbation.

In the frequency domain,

$$\widehat{\delta \phi}(\vec{r}, \omega) = \phi_0(\vec{r}) \frac{\widehat{\delta P}(\omega)}{P_0} + P_0 \widehat{\delta \psi}(\vec{r}, \omega). \quad (4.22)$$

The global amplitude is governed by the point kinetics equations and, at first order, one obtains

$$\widehat{\delta P}(\omega) = P_0 G_0(\omega) \widehat{\delta \rho}(\omega), \quad (4.23)$$

where $G_0(\omega)$ is the zero power reactor transfer function [1]

$$G_0(\omega) = \frac{1}{i\omega [\Lambda + \beta/(i\omega + \lambda)]}. \quad (4.24)$$

Substituting Eq. (4.23) into Eq. (4.22) yields

$$\widehat{\delta \phi}(\vec{r}, \omega) = \phi_0(\vec{r}) G_0(\omega) \widehat{\delta \rho}(\omega) + P_0 \widehat{\delta \psi}(\vec{r}, \omega). \quad (4.25)$$

This expression provides a useful separation between a global amplitude response and a residual shape variation. The behaviour of the point kinetic contribution is characterised by the modulus and phase of $G_0(\omega)$,

$$|G_0(\omega)|, \quad \arg G_0(\omega). \quad (4.26)$$

As illustrated in Fig. 4.1, three main frequency regions can be distinguished [1]: a low frequency region dominated by delayed neutrons, an intermediate plateau region, and a high frequency region dominated by prompt neutron dynamics. Of particular importance is the plateau region,

$$\lambda \ll \omega \ll \frac{\beta}{\Lambda}, \quad (4.27)$$

for which

$$|G_0(\omega)| \simeq \frac{1}{\beta}, \quad (4.28)$$

and the phase lag is close to zero [1]. In this regime, the global power responds almost in phase with the reactivity fluctuation and with an approximately frequency independent gain. This classification into low, plateau, and high frequency regimes will be useful later in the interpretation of the numerical results for the homogeneous benchmark.




Figure 4.1: Amplitude and phase of the zero power reactor transfer function $G_0(\omega)$ [1].

Chapter 5

Monte Carlo for Neutron Noise in MGMC

MGMC is a Monte Carlo mini-app for neutron transport developed at CEA. In its standard use, it solves multi-group steady-state transport problems by tracking individual neutron histories in arbitrary geometries and by estimating the effective multiplication factor through power iteration. In addition to this conventional criticality mode, MGMC includes dedicated algorithms for neutron noise calculations in the frequency domain. These algorithms extend the standard Monte Carlo random walk by introducing complex-valued particle weights and, in the non-perturbative case, an additional discrete coordinate associated with the harmonic index.

This chapter presents the frequency-domain Monte Carlo formulations implemented in MGMC. The aim is to illustrate how the noise equations given in Chapter 4 are interpreted and sampled stochastically in practice.

5.1 Frequency-domain Monte Carlo formulation in MGMC

The Monte Carlo noise mode of MGMC is built directly on the frequency-domain transport equations recalled in Chapter 4. In the orthodox linear formulation, the noise equation reads

$$\widehat{\mathcal{B}}_0(\omega) \widehat{\delta\varphi}(\omega) = -\widehat{\delta\mathcal{B}}(\omega) \varphi_0, \quad (5.1)$$

which is a complex-valued fixed source transport problem at each angular frequency ω .

In the non-perturbative formulation, the product between perturbation and fluctuating flux is retained. One therefore works directly with the total flux and obtains

$$\widehat{\mathcal{B}}_0(\omega) \widehat{\varphi}(\omega) + \frac{1}{2\pi} (\widehat{\delta\mathcal{B}} * \widehat{\varphi})(\omega) = 0, \quad (5.2)$$

where the convolution in frequency couples different harmonic components of the solution.

The practical difference between the two formulations is fundamental. In the linear case, each frequency can be treated independently and the problem reduces to a sequence of fixed-source simulations for each harmonic produced by the source. In the non-perturbative case, the frequency becomes a true variable of the problem: all harmonics are coupled and must be solved simultaneously. For a periodic perturbation of fundamental angular frequency ω_0 , both the perturbation and the flux are represented as discrete harmonic expansions,

$$\widehat{\delta\mathcal{B}}(\omega) = 2\pi \sum_{m \in \mathbb{Z}} \widehat{\delta\mathcal{B}}_m \delta(\omega - m\omega_0), \quad \widehat{\varphi}(\omega) = 2\pi \sum_{n \in \mathbb{Z}} \widehat{\varphi}_n \delta(\omega - n\omega_0), \quad (5.3)$$

so that Eq. (5.2) becomes an infinite coupled system for the harmonic coefficients $\widehat{\varphi}_n$.

Complex-weighted random walk

In both linear and non-perturbative noise modes, MGMC solves the frequency-domain equations by sampling neutron histories using complex weights. The underlying geometry, material layout and nuclear data are the same as in the steady state Monte Carlo calculation, and the particle trajectories are generated through the usual sequence of free flights and collisions. The effect of the frequency-domain formulation appears only through multiplicative complex factors applied to the weights and through additional reaction channels associated with the perturbation operator [14].

The key idea is therefore to preserve a random walk structure while encoding the oscillatory behaviour of the solution in the particle weight. What remains to be sampled is the action of the frequency-dependent transport operators on the neutron histories. Each history carries a complex weight $w \in \mathbb{C}$, and the required observables, such as flux harmonics, reaction rate noise, or detector responses, are estimated as complex-valued tallies averaged over many histories [9, 14].

5.1.1 Treatment of the $i\omega/v$ term

A specific feature of the frequency-domain noise equation is the presence of the term $i\omega/v$ in the operator. In MGMC, this term can be handled in two ways,

corresponding to the flight-based formulation introduced by Yamamoto [8] and the collision-based formulation proposed by Rouchon *et al.* [9].

In the flight-based scheme, the particle free path is sampled with the physical total cross section Σ_t , exactly as in ordinary transport. After a flight of length s , however, the complex weight is multiplied by the phase factor

$$\exp\left(-\frac{i\omega s}{v}\right). \quad (5.4)$$

In this approach, the frequency-dependent effect is accumulated continuously along the track.

In the collision-based scheme, the same term is represented through an additional fictitious collision channel. One introduces the modified real total cross section

$$\tilde{\Sigma}_t^{\text{cb}} = \Sigma_t + \eta \frac{\omega}{v}, \quad (5.5)$$

where $\eta > 0$ is a user-defined parameter, together with an extra collision channel of cross section

$$\Sigma_{\text{copy}} = \eta \frac{\omega}{v}. \quad (5.6)$$

Whenever such a special collision is sampled, the particle continues its history with unchanged phase space coordinates, but its weight is multiplied by

$$\frac{\eta - i}{\eta}. \quad (5.7)$$

For this reason, this reaction channel is called a ‘copy’ collision. In this way, the same frequency-dependent operator is represented through a sequence of discrete weight updates at collision points rather than through a continuous phase accumulation along flights.

Delayed neutron contributions are treated in both schemes through the frequency-dependent weight-correction factor

$$\frac{\lambda_q}{\lambda_q + i\omega} \quad (5.8)$$

for precursor family q . The flight-based and collision-based approaches are both unbiased, but they lead to different statistical convergence properties.

5.1.2 Flux estimators

In the frequency-domain noise mode, MGMC can estimate the complex flux using either a collision estimator or a track-length estimator, in analogy with the standard

transport estimators introduced in Chapter 3. Both estimators accumulate complex-valued scores, since the particle weights carry the phase information associated with the frequency-domain formulation. This estimator choice is distinct from the treatment of the $i\omega/v$ term described above. The flight-based and collision-based schemes define how the frequency-dependent streaming contribution is sampled during the random walk, whereas the collision and track-length estimators define how the resulting complex flux is tallied. In principle, the two choices can therefore be combined, leading to four scheme-estimator combinations. In practice, however, a known implementation issue had affected the combination of the flight-based scheme with the track-length estimator in MGMC. The correction of this issue is verified separately in Chapter 6, where a dedicated verification benchmark is introduced to isolate the FB+TL combination and verify the corrected implementation.

5.1.3 Linear noise mode

In linear mode, MGMC interprets Eq. (5.1) as a fixed-source transport problem. For each prescribed frequency, the code first uses the stationary eigenvalue solution to sample the source term

$$-\widehat{\delta\mathcal{B}}(\omega)\varphi_0, \quad (5.9)$$

and then estimates the corresponding complex noise flux $\widehat{\delta\varphi}(\omega)$ through a Monte Carlo simulation with complex weights.

Since the equations are decoupled in frequency, one simulation is performed for each frequency of interest. This makes the linear formulation comparatively simple: the stochastic problem is essentially a standard fixed-source transport calculation, except that the source term and the tallies are complex-valued, and that the random walk rules must take into account the modifications illustrated above. In this framework, weight cancellation (see Sec. 5.1.6) is not strictly required for the simulation, but it can still be highly beneficial in practice when different contributions of opposite sign or phase must cancel within the same tally. This is particularly relevant for localised perturbations, oscillatory responses, and detector signals, where the variance associated with destructive interference may otherwise become large.

5.1.4 Non-perturbative noise mode

In the non-perturbative case, the situation is fundamentally different from the linear one. Because the unknown flux appears on both sides of the frequency-domain equation through the convolution term, the problem cannot be treated as an externally driven fixed-source calculation at each frequency. Instead, as discussed in Section 4.2.2, the coupled harmonic system is recast as an eigenvalue problem

in an enlarged phase space including the harmonic index, and it is solved by a power-iteration scheme, in close analogy with a standard criticality calculation. In this framework, the harmonic index is treated as an additional discrete phase-space coordinate. Each particle is therefore labelled not only by its position, direction, energy and weight, but also by an integer harmonic index n associated with the coefficient $\hat{\varphi}_n$. A particle belonging to harmonic n contributes only to the tallies of that harmonic [14].

The convolution term is sampled by interpreting the harmonic components of the perturbation operator as additional non-analog reaction channels. If a particle at harmonic n undergoes a perturbation reaction associated with harmonic order m , it is transferred to harmonic

$$n' = n + m, \quad (5.10)$$

with a weight update that carries the amplitude and phase of the corresponding perturbation component [14]. In this way, the coupling among harmonics is represented directly at the level of particle histories.

The treatment of the particles generated by these perturbation reactions depends on the eigenvalue formulation used for the non-perturbative problem. The two formulations introduced in Section 4.2.2 lead to different Monte Carlo power-iteration schemes, which are described below.

Original k -scheme

The original k -scheme corresponds to Eq. (4.16). In this formulation, the stationary fission operator and the perturbation of the fission operator are placed on the right hand side of the eigenvalue problem. The particles generated by these production terms are therefore banked and used to define the source of the next generation, as in a standard Monte Carlo power iteration. By contrast, the perturbations of the total and scattering operators remain on the left hand side of Eq. (4.16). Their Monte Carlo interpretation is therefore internal to the transport process within the current generation. In practice, perturbation-induced scattering-like or removal-like events are sampled during the random walk and modify the particle population before the source renewal step.

This splitting was found to be satisfactory for the homogeneous configurations considered in Chapter 6. However, it becomes problematic for the vibrating rod benchmark at realistic vibration amplitudes.

Reformulated k' -scheme

The reformulated k' -scheme corresponds to Eq. (4.17). In this formulation, all perturbation-induced terms are moved to the right hand side of the eigenvalue

problem. The left hand side contains only the time averaged transport, removal and scattering operators. As a result, the particles are transported within each generation according to the unperturbed operator, while all perturbation-induced contributions are treated as production terms for the next generation.

From the Monte Carlo point of view, this means that perturbation-induced fission, scattering and removal contributions are all banked for the renewal of the source at the end of the generation. The corresponding source particles may carry signed or complex weights, depending on the harmonic component and on the perturbation operator from which they originate. They are then processed collectively before the next generation is started.

This formulation is better suited to the vibrating rod benchmark, because it avoids the uncontrolled growth of the particle population inside a generation. In the current implementation for the rod model of Chapter 7, cancellation is applied to every production term generated for the renewal of the source at the end of each generation. This cancellation strategy is tractable here because the benchmark has a discrete multi-group, two-direction structure. The situation would be more difficult in a more general transport setting, where extending cancellation to all production channels would require handling a broader and possibly continuous phase space.

The reformulated k' -scheme is the non-perturbative Monte Carlo algorithm used for the one-dimensional vibrating rod calculations presented in Chapter 7.

5.1.5 Symmetry and frequency mirror

The harmonic components satisfy the usual symmetry relations

$$\widehat{\delta\mathcal{B}}_{-n} = \widehat{\delta\mathcal{B}}_n^*, \quad \widehat{\varphi}_{-n} = \widehat{\varphi}_n^*. \quad (5.11)$$

where the asterisk denotes the complex conjugation. This symmetry can be exploited to reduce the computational cost by restricting the simulation to non-negative harmonics only [14].

MGMC implements this reduction through a frequency mirror at $n = 0$ [14]. Whenever a perturbation reaction would send a particle to a negative harmonic, the particle is reflected to the corresponding positive harmonic and its weight is conjugated consistently with the symmetry relation above. Special care is taken for the zero harmonic, which corresponds to the time-averaged component of the solution: as both the particle and its image land on the mirror, the particle weight is set to twice its real part. This also ensures that tallies associated with $n = 0$ remain purely real.

This construction avoids an explicit simulation of negative harmonics, while preserving the exact spectral content of the time-domain solution.

5.1.6 Weight cancellation

The solution strategy differs significantly between the linear and non-perturbative modes. In linear mode, as noted above, the problem is a fixed-source one, and no eigenvalue iteration is needed. In non-perturbative mode, by contrast, the coupled harmonic system is solved by power iteration in the enlarged phase space [14].

A major difficulty then arises from the use of signed or complex valued weights. In such situations, the dominant eigenpair of the naive Monte Carlo power iteration is not the physical one, because the signed measure associated with the transport operator favours an unphysical mode with the largest modulus [29]. For this reason, weight cancellation is an essential ingredient of the non-perturbative algorithm.

Weight cancellation consists in combining nearby contributions of opposite sign or phase so as to suppress unphysical growth associated with signed or complex weights [29]. In the present non-perturbative formulation, cancellation is performed separately for each harmonic, since harmonics represent discrete locations in the extended phase space [14]. Besides reducing the variance associated with large cancellations between opposite-weight contributions, this procedure is crucially required to suppress the non-physical dominant eigenpair and to restore convergence of power iteration towards the physical eigenpair [29, 14, 10]. Weight cancellation must be performed collectively on all particles; therefore, it is applied at the end of each generation, before population normalization.

Although weight cancellation is not strictly required in linear fixed-source calculations, it can still be used as a variance-reduction technique. In that case, in order to perform weight cancellation collectively, particle histories must be stopped at suitable rendez-vous points. In practice, the transport routine is reorganised into fission generations, similarly to criticality calculations. At the end of each generation, particles undergo cancellation before resuming transport. Unlike in eigenvalue calculations, no population normalisation is applied after cancellation.

Chapter 6

Homogeneous Medium

The homogeneous medium configuration considered in this chapter provides a minimal test-bed for the verification of frequency-domain neutron noise calculations in MGMC. Since spatial dependence is suppressed by construction, the noise problem reduces to a point-kinetics-like system. The frequency response is therefore governed entirely by the energy structure, the delayed neutron parameters, and the time dependence of the perturbation.

This simplified setting is first used to verify a specific implementation feature of the Monte Carlo solver, namely the track-length estimator combined with the flight-based treatment of the $i\omega/v$ term. In order to make an implementation error affecting this estimator clearly observable, a dedicated scattering-free high-frequency case is introduced for this purpose. The remainder of the chapter then focuses on the comparison between the linear and non-perturbative (NP) frequency-domain formulations.

6.1 Verification of the flight-based scheme with track-length estimator

As discussed in Chapter 5, the noise mode of MGMC makes it possible to estimate the complex frequency-domain flux either with a collision estimator or with a track-length estimator. A previous implementation error affecting the flight-based scheme with track-length estimator had been identified, but its effect remained difficult to observe in standard homogeneous simulations. For this reason, a dedicated verification test was designed in which the FB+TL combination could be isolated and compared against an analytical reference solution. The objective of the present section is therefore to verify the corrected implementation of this estimator in the

MGMC mini-app.

In the flight-based formulation introduced by Yamamoto [8], the contribution of the $i\omega/v$ term is accounted for during the particle flight by assigning a complex phase factor to the neutron weight. As a consequence, the corresponding track-length estimator must account for the variation of the complex weight along the flight segment. Consider a particle in group g undergoing the j -th free flight, with distance to collision s_j and speed $v_{g,j}$. Let $w_j(s)$ denote the complex weight as a function of the distance travelled along the segment, and let w_{j-1} be the weight at the beginning of the flight, i.e. immediately after the previous collision or source event. In the flight-based treatment of the $i\omega/v$ term, the weight varies continuously along the trajectory. Therefore, the contribution of this segment to the frequency-domain flux is given by [8, 9]

$$\text{TL}_j(\omega) = \int_0^{s_j} w_j(s) ds = \int_0^{s_j} w_{j-1} \exp\left(-\frac{i\omega}{v_{g,j}}s\right) ds. \quad (6.1)$$

Eq. (6.1) can be integrated, yielding

$$\text{TL}_j(\omega) = w_{j-1} \frac{v_{g,j}}{i\omega} \left[1 - \exp\left(-\frac{i\omega}{v_{g,j}}s_j\right) \right]. \quad (6.2)$$

In practice, the value yielded by Eq. (6.2) can only be distinguished from the usual formula $\text{TL}_j(\omega) = w_j l_j$ if the weight varies sufficiently along flight segments. Therefore, in order to provide a relevant unit test for the implementation of this estimator, a dedicated verification case had to be designed.

6.1.1 Scattering-free infinite-medium test problem

The verification is performed in a one-group homogeneous infinite medium, represented in MGMC by a single reflective cell. The configuration is chosen without scattering in order to generate long uninterrupted flights. The macroscopic cross sections and kinetic parameters are listed in Table 6.1. Prompt and delayed neutron multiplicities are the same as in Table 6.4.

The perturbation is a spatially uniform periodic modulation of the fission cross section,

$$\Sigma_f(t) = \Sigma_{f,0} [1 + \varepsilon_f \cos(\omega_0 t)], \quad \varepsilon_f = 5 \times 10^{-3}. \quad (6.3)$$

Only the fundamental harmonic is retained, and the frequency is fixed at $f = 10^4$ Hz, corresponding to $\omega_0 = 2\pi f = 6.283 \times 10^4 \text{ s}^{-1}$.

For this test, the linear formulation is used and no non-perturbative eigenvalue scheme is involved. The MGMC input is configured with either the flight-based

or the collision-based treatment of the $i\omega/v$ term. For each choice, both the collision estimator and the track-length estimator are activated, leading to four combinations:

$$\text{FB+TL}, \quad \text{FB+CL}, \quad \text{CB+TL}, \quad \text{CB+CL}.$$

Analytical reference values for the real and imaginary parts of the flux are obtained by solving the corresponding one-group frequency-domain equations for the same configuration. The Monte Carlo estimates are compared with these reference values by means of a two-sided Gaussian test. For a generic quantity q , with Monte Carlo estimate $\hat{q}$, estimated standard error $\sigma_{\hat{q}}$, and analytical reference q_{ref} , the z -score is defined as

$$z = \frac{\hat{q} - q_{\text{ref}}}{\sigma_{\hat{q}}}, \quad (6.4)$$

and the corresponding two-sided p -value is

$$p = 2[1 - \Phi(|z|)], \quad (6.5)$$

where Φ is the cumulative distribution function of the standard normal distribution [30]. In this work, values $p < 0.05$ are taken to indicate rejection of the null hypothesis and are therefore reported as statistically significant discrepancies.

For each of the four combinations, all MGMC simulations were run with 10^4 particles per batch over 500 batches.

Table 6.1: Macroscopic cross sections and kinetic parameters for the one-group, scattering-free validation case.

Quantity	Symbol	Value	Unit
Total macroscopic cross section	Σ_t	1.0	cm^{-1}
Scattering cross section	Σ_s	0.0	cm^{-1}
Absorption cross section	Σ_a	1.0	cm^{-1}
Fission cross section	Σ_f	0.4	cm^{-1}
Delayed fraction	β	7×10^{-3}	–
Precursor decay constant	λ	0.08	s^{-1}
Neutron speed	v	3.0×10^5	cm s^{-1}

6.1.2 Results before and after the correction

Before correction, the scattering-free high-frequency test reveals a clear inconsistency affecting the flight-based scheme with track-length estimator. Table 6.2 compares

the analytical reference flux with the Monte Carlo estimates obtained from the four scheme-estimator combinations.

Table 6.2: Scattering-free test at $f = 10^4$ Hz: comparison between analytical reference and Monte Carlo fluxes before correcting the implementation of the track-length estimator.

Scheme	Quantity	Value	p -value
Reference	$\Re\{\hat{\varphi}\}$	2.49×10^{-3}	–
Reference	$\Im\{\hat{\varphi}\}$	-7.44×10^{-2}	–
FB+TL	$\Re\{\hat{\varphi}\}$	1.70×10^{-2}	0
FB+TL	$\Im\{\hat{\varphi}\}$	-7.39×10^{-2}	0
FB+CL	$\Re\{\hat{\varphi}\}$	2.50×10^{-3}	0.527
FB+CL	$\Im\{\hat{\varphi}\}$	-7.44×10^{-2}	0.997
CB+TL	$\Re\{\hat{\varphi}\}$	2.51×10^{-3}	0.692
CB+TL	$\Im\{\hat{\varphi}\}$	-7.43×10^{-2}	0.247
CB+CL	$\Re\{\hat{\varphi}\}$	2.52×10^{-3}	0.631
CB+CL	$\Im\{\hat{\varphi}\}$	-7.43×10^{-2}	0.270

The FB+TL combination exhibits a statistically significant discrepancy in both the real and imaginary parts of the flux, with essentially zero p -values. By contrast, the collision estimator in the flight-based scheme and both estimators in the collision-based scheme remain statistically compatible with the analytical reference, with $p > 0.05$.

After correcting the implementation of the flight-based track-length estimator in MGMC, the same test was repeated. Table 6.3 reports the analytical reference and the corrected FB+TL estimate.

Table 6.3: Scattering-free test at $f = 10^4$ Hz: analytical reference and corrected track-length estimator.

Scheme	Quantity	Value	p -value
Reference	$\Re\{\hat{\varphi}\}$	2.49×10^{-3}	–
Reference	$\Im\{\hat{\varphi}\}$	-7.44×10^{-2}	–
FB+TL corrected	$\Re\{\hat{\varphi}\}$	2.50×10^{-3}	0.383
FB+TL corrected	$\Im\{\hat{\varphi}\}$	-7.44×10^{-2}	0.366

After correction, the FB+TL estimate recovers agreement with the analytical reference within statistical uncertainty, with both p -values well above the 0.05

threshold. The other scheme-estimator combinations are unchanged and remain compatible with the reference. This verification therefore confirms the corrected implementation of the flight-based scheme with track-length estimator in MGMC.

6.2 Linear vs non-perturbative comparison

The homogeneous medium configuration considered in this section provides a minimal test-bed for comparing the linear and non-perturbative (NP) frequency domain formulations implemented in MGMC.

Three families of simulations are considered, all using one effective delayed neutron family:

1. one energy group, single-harmonic perturbation;
2. one energy group, two-harmonic perturbation;
3. two energy groups, single-harmonic perturbation.

For each family, six frequencies are investigated,

$$f \in \{10^{-1}, 1, 10, 10^2, 10^3, 10^4\} \text{ Hz},$$

covering the plateau region and the high-frequency regime, and approaching the low-frequency transition.

6.2.1 Problem definition and input data

The geometry consists of a single cubic cell with reflective boundary conditions, representing an infinite homogeneous medium from the transport point of view. The perturbation is a spatially uniform modulation of the fission cross section. Scattering is assumed isotropic in angle. Delayed neutrons are represented by one effective precursor group with decay constant $\lambda = 0.08 \text{ s}^{-1}$.

One-group configuration

In the one-group configuration, the homogeneous medium is characterised by the macroscopic cross sections and kinetic parameters reported in Table 6.4. The time-averaged system is adjusted to criticality, and the reference solution $\hat{\varphi}_0$ is obtained from a standard k -eigenvalue calculation.

Table 6.4: Macroscopic cross sections and kinetic parameters for the one-group homogeneous configuration.

Quantity	Symbol	Value	Unit
Total macroscopic cross section	Σ_t	2.0	cm^{-1}
Scattering cross section	Σ_s	1.0	cm^{-1}
Absorption cross section	Σ_a	1.0	cm^{-1}
Fission cross section	Σ_f	0.4	cm^{-1}
Prompt neutron multiplicity	ν_p	2.4825	–
Delayed neutron multiplicity	ν_d	0.0175	–
Neutron speed	v	3.0×10^5	cm s^{-1}

Two-group configuration

In the two-group configuration, the homogeneous medium is described by the multi-group cross sections, scattering matrix, neutron speeds, and kinetic parameters listed in Table 6.5. Group 0 denotes the fast group and group 1 the thermal group. Fission neutrons are emitted in the fast group only, so that $\chi = (1,0)$ for both prompt and delayed neutrons. As in the one-group case, the time-averaged system is adjusted to criticality.

Table 6.5: Macroscopic cross sections, scattering matrix, neutron speeds, and kinetic parameters for the two-group homogeneous configuration.

Quantity	Symbol	Group 0	Group 1	Unit
Total macroscopic cross section	$\Sigma_{t,g}$	0.285	2.1	cm^{-1}
Absorption cross section	$\Sigma_{a,g}$	5×10^{-3}	0.1	cm^{-1}
Fission cross section	$\Sigma_{f,g}$	8×10^{-4}	4.4×10^{-2}	cm^{-1}
Scattering cross section $0 \rightarrow 0$	$\Sigma_{s,0 \rightarrow 0}$	0.25	–	cm^{-1}
Scattering cross section $0 \rightarrow 1$	$\Sigma_{s,0 \rightarrow 1}$	–	0.03	cm^{-1}
Scattering cross section $1 \rightarrow 0$	$\Sigma_{s,1 \rightarrow 0}$	0.00	–	cm^{-1}
Scattering cross section $1 \rightarrow 1$	$\Sigma_{s,1 \rightarrow 1}$	–	2.0	cm^{-1}
Neutron speed	v_g	2.8×10^7	3.0×10^5	cm s^{-1}
Prompt neutron multiplicity	$\nu_{p,g}$	2.483025	2.483025	–
Delayed neutron multiplicity	$\nu_{d,g}$	0.016975	0.016975	–
Fission spectrum	χ_g	1.0	0.0	–

Time-dependent perturbation

In all cases, the perturbation acts as a spatially uniform modulation of the fission cross section,

$$\Sigma_f(t) = \Sigma_{f,0} [1 + \varepsilon_f \cos(\omega_0 t)], \quad \varepsilon_f = 5 \times 10^{-3}. \quad (6.6)$$

In the present chapter, the perturbation is described in terms of the harmonic components retained in the calculation. In the single-harmonic cases, only the fundamental angular frequency $\omega_0 = 2\pi f$ is included. In the two-harmonic case, the perturbation contains both ω_0 and $2\omega_0$.

Frequency range considered

The analysis is restricted to frequencies

$$f \geq 10^{-1} \text{ Hz},$$

that is, above the low-frequency transition. Very low frequencies are not considered because, in the proof of concept Monte Carlo study of the same infinite-medium benchmark, the dominance ratio of the NP transition operator was found to be extremely close to unity, which led to very long source convergence times [14]. When the coupled harmonic system is truncated, the corresponding deterministic transition matrix may even develop complex conjugate eigenvalues with modulus larger than one, so that power iteration does not converge [14]. Related numerical difficulties at very low frequencies have also been reported for frequency-domain Monte Carlo in the linear framework [5].

For this reason, the present validation focuses on the plateau and high-frequency regimes, where both deterministic and Monte Carlo NP calculations exhibit robust convergence.

6.2.2 Numerical computation settings

All MGMC calculations for the homogeneous-medium benchmarks were performed with the flight-based noise scheme, using delta-tracking and the collision estimator. In the linear formulation, which is treated as a fixed-source problem, the one-group single-harmonic and one-group two-harmonic cases were run with 10^4 particles per batch over 500 batches. The two-group single-harmonic case was run with 500 particles per batch over 1500 batches. In the non-perturbative formulation, the coupled harmonic system was solved as an eigenvalue problem by power iteration in the extended phase space. Here, a generation denotes one iteration of the eigenvalue scheme, whereas the inactive generations are discarded in order to allow source

convergence before tally accumulation. For these calculations, the original k -scheme was used. The one-group single-harmonic and one-group two-harmonic cases were run with 10^5 particles per generation over 4000 generations, of which the first 2000 were inactive. The two-group single-harmonic case was run with 5×10^3 particles per generation, again over 4000 generations, with 2000 inactive generations. Each non-perturbative result was obtained from 10 independent simulations and the non-perturbative harmonic hierarchy was truncated at a maximum harmonic index $h_{\max} = 50$.

Reference solutions for the homogeneous 0D benchmarks were obtained with dedicated in-house scripts implementing the corresponding frequency-domain equations. In the linear formulation, analytical reference values for the real and imaginary parts of the flux were obtained by solving the corresponding frequency-domain algebraic equations for the same homogeneous configuration. In the non-perturbative formulation, the coupled harmonic hierarchy was truncated at a maximum harmonic index $h_{\max} = 50$ and rewritten as a finite-dimensional eigenvalue problem. The dominant eigenpair was then computed by power iteration, and the harmonic components of the flux were extracted from the converged eigenvector.

In the figures of this chapter, the following shorthand notation is used: *Linear* denotes the linear formulation, *NP* the non-perturbative formulation, *MC* Monte Carlo results, and *Det* deterministic reference solutions, obtained analytically for the linear formulation and by deterministic power iteration for the non-perturbative one. Unless otherwise stated, error bars represent one standard deviation and are generally smaller than the marker size.

6.2.3 One-group, single-harmonic perturbation

The first family of tests considers one energy group and a perturbation containing only the fundamental harmonic $\omega_0 = 2\pi f$. For each frequency, the amplitude and phase of the first harmonic of the normalised flux are compared between the linear and NP Monte Carlo calculations and, for this case, also with deterministic reference solutions.

Figure 6.1 shows the amplitude and phase of the first harmonic as a function of frequency. The left panel reports $|\hat{\varphi}_1|$, while the right panel shows $\arg(\hat{\varphi}_1)$. Over the whole frequency range considered, the Monte Carlo estimates closely follow the deterministic curves in both amplitude and phase. This provides a first verification of the implementation for the simplest benchmark configuration. In the plateau region, the NP first-harmonic amplitude is systematically smaller than the linear one, whereas in the high frequency regime the two formulations become practically indistinguishable within the statistical uncertainty.




Figure 6.1: One-group, single-harmonic case: amplitude (left) and phase (right) of the normalised first harmonic as a function of frequency.

The observed frequency response exhibits the expected structure: an intermediate plateau followed by a high frequency decay. As discussed in Section 4.2.3, these regimes are separated by two characteristic frequencies derived from the one-group point-kinetics model. The first, associated with the precursor decay constant, is

$$f_{\text{low}} = \frac{\lambda}{2\pi} = 0.01273 \text{ Hz}, \quad (6.7)$$

whereas the second, associated with prompt neutron dynamics modulated by the delayed fraction, is

$$f_{\text{high}} = \frac{\beta \nu \Sigma_f v}{2\pi} = 334.2 \text{ Hz}. \quad (6.8)$$

In Fig. 6.1, these two frequencies are indicated by the vertical dashed black line (f_{low}) and red line (f_{high}). Their positions are consistent with the observed change of slope in the amplitude and with the onset of the phase roll off, which provides an additional consistency check between the numerical results and the analytical estimates.

To clarify the origin of the differences between the linear and NP results, two representative frequencies are selected: one in the plateau region ($f = 10$ Hz) and one in the high frequency regime ($f = 10^3$ Hz). For each of them, the harmonic content of the NP solution is analysed through the modulus and phase of the Fourier coefficients as a function of the harmonic index h ,

$$|\hat{\varphi}_h|, \quad \arg(\hat{\varphi}_h), \quad h = 0, 1, \dots, 4.$$

Figures 6.2 and 6.3 report the corresponding results. Since the linear formulation retains only the driving harmonic, all higher harmonics are zero by construction in the linear case, whereas the NP solution exhibits a non trivial harmonic spectrum.




Figure 6.2: One-group, single-harmonic case: modulus (left) and phase (right) of $\hat{\varphi}_h$ as a function of harmonic index h at a representative plateau frequency ($f = 10$ Hz).




Figure 6.3: One-group, single-harmonic case: modulus (left) and phase (right) of $\hat{\varphi}_h$ as a function of harmonic index h at a representative high frequency ($f = 10^3$ Hz).

At the plateau frequency, the NP solution exhibits a clear hierarchy of non-zero higher harmonics. Their modulus decreases with h , but the contributions at $h = 2$ and $h = 3$ remain non negligible, and the corresponding phases display a structured dependence on the harmonic index. These additional components arise from the harmonic coupling retained in the NP formulation. Their presence explains why

the NP and linear first-harmonic responses do not coincide exactly in the plateau region: part of the spectral content is redistributed among higher harmonics in the NP solution, whereas the linear formulation constrains the response to the driving component only.

In the high-frequency regime, the situation changes markedly. The NP higher harmonics are strongly damped and their modulus rapidly decreases towards zero. The influence of harmonic coupling on the fundamental component then becomes weak, and the NP first-harmonic response becomes essentially indistinguishable from the linear one. This is consistent with the convergence of the linear and NP curves already observed in Fig. 6.1.

6.2.4 One-group, two-harmonic perturbation

The second family of tests considers the same one-group homogeneous medium, but now the perturbation contains the first two harmonics of the base frequency. The six frequencies

$$f \in \{10^{-1}, 1, 10, 10^2, 10^3, 10^4\} \text{ Hz}$$

are again explored.

Figures 6.4 and 6.5 show the amplitude and phase of the first and second harmonics as functions of frequency. For both $h = 1$ and $h = 2$, the same qualitative behaviour as in the previous case is observed: an intermediate plateau followed by a high-frequency decay. In the plateau region, the NP and linear responses differ systematically, with the NP amplitudes larger in magnitude. For $f > f_{\text{high}}$, the two formulations become nearly indistinguishable for both harmonics.




Figure 6.4: One-group, two-harmonic case: amplitude (left) and phase (right) of the first harmonic $\hat{\varphi}_1$ as a function of frequency.




Figure 6.5: One-group, two-harmonic case: amplitude (left) and phase (right) of the second harmonic $\hat{\varphi}_2$ as a function of frequency.

To characterise the role of harmonic coupling more explicitly, Figures. 6.6 and 6.7 report the modulus and phase of $\hat{\varphi}_h$ as functions of h at a representative plateau frequency ($f = 10$ Hz) and at a representative high frequency ($f = 10^3$ Hz), respectively.




Figure 6.6: One-group, two-harmonic case: modulus (left) and phase (right) of $\hat{\varphi}_h$ as a function of harmonic index h at a representative plateau frequency ($f = 10$ Hz).




Figure 6.7: One-group, two-harmonic case: modulus (left) and phase (right) of $\hat{\varphi}_h$ as a function of harmonic index h at a representative high frequency ($f = 10^3$ Hz).

In the plateau regime, the NP solution exhibits a non-trivial harmonic hierarchy. Beyond the two driving components $h = 1$ and $h = 2$, additional harmonics with $h \geq 3$ are generated, with decreasing but still noticeable amplitude. Their phases also show a structured dependence on h . In the linear formulation, by contrast, only the explicitly driving harmonics are present and all the others vanish identically. The richer spectral content of the NP solution explains the systematic differences between the linear and NP first and second harmonic responses observed in the plateau region.

At high frequency, the NP higher harmonics are strongly damped and collapse towards zero. The spectrum is then effectively limited to the driving components, and the linear and NP results coincide within statistical uncertainty. This confirms, also in the two-harmonic case, that the impact of NP coupling is strongest in the plateau regime and becomes negligible in the high-frequency limit.

6.2.5 Two-group, single-harmonic perturbation

The third family of simulations considers the two-group homogeneous medium defined in Section 6.2.1, with a perturbation containing only the fundamental harmonic $\omega_0 = 2\pi f$. The same six frequencies are investigated.

Figures 6.8 and 6.9 show the amplitude and phase of the first harmonic of the group-wise flux as functions of frequency for the fast and thermal groups, respectively. For both energy groups, the qualitative behaviour mirrors the one-group case:

an approximately flat plateau for $f_{\text{low}} < f < f_{\text{high}}$, followed by a decay of the amplitude at high frequency. In the plateau region, the NP amplitudes are smaller than the linear ones, while for $f > f_{\text{high}}$ the two formulations become essentially indistinguishable within statistical uncertainty.




Figure 6.8: Two-group, single-harmonic case: amplitude (left) and phase (right) of the first-harmonic fast-group flux $\hat{\varphi}_{1,0}$ as a function of frequency.




Figure 6.9: Two-group, single-harmonic case: amplitude (left) and phase (right) of the first-harmonic thermal-group flux $\hat{\varphi}_{1,1}$ as a function of frequency.

Figures 6.10 and 6.11 show the modulus and phase of $\hat{\varphi}_{h,1}$ as functions of the harmonic index h at a plateau frequency ($f = 10$ Hz) and at a high frequency ($f = 10^2$ Hz), respectively. For brevity, only the thermal group is reported, since the fast group exhibits the same qualitative behaviour.




Figure 6.10: Two-group, single-harmonic case: modulus (left) and phase (right) of the thermal-group coefficients $\widehat{\varphi}_{h,1}$ as a function of harmonic index h at a plateau frequency ($f = 10$ Hz).




Figure 6.11: Two-group, single-harmonic case: modulus (left) and phase (right) of the thermal-group coefficients $\widehat{\varphi}_{h,1}$ as a function of harmonic index h at a high frequency ($f = 10^2$ Hz).

In the plateau regime, the NP solution again exhibits a non-trivial spectrum of higher harmonics, whereas the linear solution contains only the fundamental harmonic. At high frequency, the NP higher-order harmonics are strongly damped and the NP and linear spectra become practically identical. The same qualitative behaviour is observed for the fast group and is not shown for brevity.

Overall, the two-group benchmark confirms that the conclusions drawn from the one-group analysis remain valid once a simple energy structure is introduced. It should be noted, however, that non-trivial effects cannot be excluded in a continuous-energy treatment, which is beyond the scope of the present work. Non-linear harmonic coupling modifies the response in the plateau regime, whereas at high frequency the orthodox and non-perturbative formulations converge towards the same first-harmonic behaviour.

Chapter 7

One-dimensional Vibrating Rod

As recalled in Chapter 4, mechanical vibrations of core components represent a particularly demanding class of neutron noise problems. In such cases, the perturbation acts on interfaces or geometrical boundaries, and the induced response may become locally comparable in magnitude to the stationary flux. Deterministic non-perturbative analyses of one-dimensional rod models have shown that, for realistic vibration amplitudes, the orthodox linear formulation may fail to reproduce the full noise response accurately: the differences between linear and non-perturbative solutions remain limited for the zero and first harmonic, but become significant for higher order harmonics of the neutron field [13, 31]. This motivates the use of the non-perturbative formulation introduced in Section 4.2.2.

On the Monte Carlo side, Chapter 5 presented the general frequency-domain implementation in MGMC for both the linear and non-perturbative formulations. A first non-perturbative eigenvalue algorithm, based on the original k -scheme introduced in Section 4.2.2, was shown to behave satisfactorily for infinite-medium configurations and for moderately perturbed systems. However, when this algorithm was applied to the one-dimensional vibrating rod benchmark with realistic vibration amplitudes, the simulations exhibited very poor convergence properties. In practice, converged results could only be obtained by reducing the vibration amplitude to unrealistically small values. For amplitudes of practical interest, the perturbed system became supercritical within each generation, leading to an effectively non convergent simulation.

The origin of these difficulties can be traced back to the splitting used in the original k -scheme. In that formulation, perturbation-induced fission events contribute to

the source of the next generation, whereas the contributions associated with perturbations of scattering and total cross sections are sampled within the current generation. For weak perturbations, this splitting remains manageable. When the vibration amplitude increases, however, the balance within a generation may become unfavourable: the perturbation-induced production generated at collisions can exceed the corresponding absorption and removal before the next source-renewal step. The resulting growth of the particle population makes the simulation extremely inefficient and prevents convergence. Weight cancellation could in principle help to control this growth, but applying it inside each generation would be computationally too expensive in the present implementation.

In view of these limitations, the reformulated k' -scheme introduced in Section 4.2.2 is adopted and validated in this work. The chapter is organised as follows. First, the vibrating rod benchmark and the associated deterministic and Monte Carlo simulation settings are presented. Second, the resulting harmonic flux profiles are compared for the linear and non-perturbative formulations. Finally, a parametric study is used to clarify the numerical settings required for reliable non-perturbative simulations and a Figure of Merit analysis is performed to assess the computational efficiency of the method.

7.1 Rod model: problem definition and numerical set up

The test case considered in this chapter is the one-dimensional rod vibration model introduced in Ref. [13]. Neutrons are confined to move along a one-dimensional domain of length $L = 10$ cm, with two material regions: a central fuel pin of nominal width $d = 6$ cm, surrounded on both sides by a moderating region. In the reference configuration, the two fuel moderator interfaces are located at $x = 2$ cm and $x = 8$ cm. Both interfaces undergo a sinusoidal vibration at frequency $f = 0.1$ Hz and amplitude $\varepsilon = 0.5$ cm, moving in phase with one another. The fuel pin therefore performs a rigid oscillatory translation within the moderator, while its width remains constant. A schematic representation of the geometry and of the vibrating interfaces is given in Fig. 7.1.

The model assumes single-speed transport and involves a single effective delayed neutron precursor family. The rod model restricts particle motion to two discrete directions, after an isotropic scattering event, either direction is selected with equal probability. The geometric and nuclear input data used in this work are summarised in Table 7.1; these values are consistent with the reference case of Ref. [13].




Figure 7.1: Schematic representation of the one-dimensional vibrating rod model.

Table 7.1: Input data for the one-dimensional vibrating rod model.

Quantity	Symbol	Value	Unit
Domain length	L	10	cm
Fuel nominal width	d	6	cm
Vibration amplitude (per interface)	ε	0.5	cm
Vibration frequency	f	0.1	Hz
Moderator scattering cross section	Σ_s^{mod}	3.0	cm^{-1}
Moderator capture cross section	Σ_c^{mod}	0.2	cm^{-1}
Fuel scattering cross section	Σ_s^{fuel}	1.0	cm^{-1}
Fuel capture cross section	Σ_c^{fuel}	1.2	cm^{-1}
Fuel fission cross section	Σ_f^{fuel}	0.98	cm^{-1}
Average fission multiplicity	ν	2.2711	–
Delayed neutron fraction	β	7×10^{-3}	–
Precursor decay constant	λ	8×10^{-2}	s^{-1}
Neutron speed	v	2.2×10^5	cm s^{-1}

The reformulated k' -scheme described in Section 4.2.2 is used to compute simultaneously the harmonics of the Fourier transformed flux $\hat{\varphi}_h(x)$ at integer multiples $h\omega_0$ of the fundamental angular frequency $\omega_0 = 2\pi f$. The harmonic index h is treated as an additional discrete phase space coordinate, and the coupled hierarchy resulting from Eq. (4.17) is truncated at a maximum harmonic h_{max} , chosen large enough to ensure that the harmonics of interest are accurately represented. The same geometry and material data are used in the linear Monte Carlo noise calculations,

where each frequency is treated independently, and in the deterministic reference solver.

7.2 Numerical computation settings

All MGMC calculations for the vibrating-rod benchmark were performed with the flight-based noise scheme, using delta-tracking and the collision estimator. In the linear formulation, each frequency was treated independently as a fixed-source problem. The corresponding simulations were run with 10^6 particles per batch over 10^4 batches, with the first 1500 inactive cycles discarded before tally accumulation. In the non-perturbative formulation, the reformulated eigenvalue problem of Eq. (4.17) was solved by power iteration in the extended phase space, using the k' -scheme. The corresponding simulations were run with 10^5 particles per generation over 3×10^4 generations, of which the first 1.5×10^4 were inactive. Each non-perturbative result was obtained from 20 independent runs, and the coupled harmonic hierarchy was truncated at $h_{\max} = 7$.

For the deterministic calculations, the one-dimensional transport equation associated with the vibrating rod model is discretised by finite differences. The generalised eigenvalue problem corresponding to Eq. (4.17) is converted implicitly into an ordinary eigenvalue problem by solving the underlying linear systems with the iterative Krylov solver GMRES [32]. The eigenpairs are then computed with the Implicitly Restarted Arnoldi Method as implemented in ARPACK [33]. This choice is well suited to the present problem because the discretised operator is large, sparse, and non symmetric, so that direct dense eigensolvers would be unnecessarily expensive.

The comparison between linear and non-perturbative Monte Carlo noise simulations and the deterministic solutions is carried out for the first few harmonics of the noise field. The analysis focuses on the spatial profiles of the normalised flux at harmonics $h = 0, 1, 2$, and 3 , and on the differences between linear and non-perturbative formulations in terms of both real and imaginary parts. The corresponding results are presented in the next section.

In addition to the representative frequency reported in detail below, the same one-dimensional benchmark was also analysed at $f = 10$ Hz and $f = 100$ Hz, which both lie in the plateau region. These additional calculations led to the same qualitative conclusions regarding the comparison between the linear and non-perturbative formulations. For brevity, they are not reported here, since they do not provide qualitatively new information beyond the case shown in the following.

7.3 Results for the vibrating rod benchmark

This section compares the reformulated (i.e. k') non-perturbative Monte Carlo algorithm with both the linear Monte Carlo formulation and the deterministic reference solutions. All harmonics are plotted with a consistent normalisation of the flux time-average. For each harmonic h , the upper part of the figure shows the normalised Fourier coefficient $\hat{\varphi}_h(x)$, while the lower part reports the difference

$$\Delta\hat{\varphi}_h(x) = \hat{\varphi}_h^{\text{MC}}(x) - \hat{\varphi}_h^{\text{det}}(x), \quad (7.1)$$

separately for the linear and non-perturbative formulations. In all plots, red curves correspond to deterministic solutions and black curves to Monte Carlo solutions; dotted lines denote the linear formulation and solid lines the non-perturbative one. The shaded bands indicate the estimated one-sigma statistical uncertainty of the Monte Carlo results.

7.3.1 Zero-frequency component

Figure 7.2 shows the real part of the zero frequency component $\hat{\varphi}_0(x)$, which coincides with the time-averaged flux. As expected, the imaginary part is identically zero and is therefore not reported. The deterministic linear, deterministic non-perturbative, linear Monte Carlo, and non-perturbative Monte Carlo curves are essentially indistinguishable at the scale of the figure. The lower panel confirms that the Monte Carlo estimates agree with the deterministic reference within a few 10^{-4} in absolute value over the whole domain.




Figure 7.2: Real part of the zero frequency component $\hat{\varphi}_0(x)$ (top) and corresponding difference $\Re[\Delta\hat{\varphi}_0(x)]$ between Monte Carlo and deterministic solutions (bottom).

This result shows that the reformulated non-perturbative algorithm reproduces the zero frequency component correctly. It confirms that the present benchmark does

not reveal major linear versus non-perturbative differences on the zero frequency component.

7.3.2 First harmonic

Figure 7.3 reports the real and imaginary parts of the first harmonic $\hat{\varphi}_1(x)$, shown respectively in the left and right panels. The imaginary part is the dominant contribution, with amplitudes about two orders of magnitude larger than those of the real part. For both components, the deterministic linear and non-perturbative solutions remain close to each other, and the Monte Carlo results follow the deterministic curves over the whole domain.

The difference plots indicate deviations at the level of a few per cent of the peak amplitude, generally within the one-sigma bands and without visible systematic drift. This confirms the conclusions of References [13, 31]: for realistic vibration amplitudes, the first harmonic is well captured by the linear formulation, and the effect of nonlinear coupling remains modest. From the Monte Carlo point of view, the agreement also shows that the reformulated eigenvalue problem (4.17) and its stochastic implementation reproduce the first harmonic correctly while preserving the accuracy of the linear algorithm.




Figure 7.3: First harmonic $\hat{\varphi}_1(x)$ of the vibrating rod model: real part (left) and imaginary part (right). Each panel shows the normalised harmonic (top) and the difference between Monte Carlo and deterministic solutions (bottom).

7.3.3 Second harmonic

A qualitatively different picture emerges for the higher order harmonics. Figure 7.4 displays the second harmonic $\hat{\varphi}_2(x)$, with the real and imaginary parts shown in the left and right panels, respectively. Here the real part is dominant. The deterministic and Monte Carlo non-perturbative solutions exhibit a relatively small, almost flat

profile in the central fuel region, with only variations near the vibrating interfaces at $x = 2$ cm and $x = 8$ cm. By contrast, the linear solutions predict a much larger, almost parabolic variation across the whole domain, with an amplitude more than one order of magnitude larger than that of the non-perturbative solutions.

The difference plots make the situation particularly clear. The non-perturbative Monte Carlo results remain in very good agreement with the deterministic non-perturbative reference. The linear Monte Carlo calculations are likewise in good agreement with the corresponding deterministic linear solution, but that linear reference profile is itself strongly different from the non-perturbative one. This is therefore the first harmonic for which the inadequacy of the linear formulation becomes relevant in the present benchmark.




Figure 7.4: Second harmonic $\hat{\varphi}_2(x)$ of the vibrating rod model: real part (left) and imaginary part (right). Each panel shows the normalised harmonic (top) and the difference between Monte Carlo and deterministic solutions (bottom).

7.3.4 Third harmonic

The third harmonic, shown in Fig. 7.5, further highlights the nonlinear nature of vibration induced noise. The imaginary part is the dominant component, while the real part remains comparatively small. The deterministic non-perturbative solution exhibits a response strongly localised around the vibrating interfaces and changing sign across them. The linear solution reproduces part of the qualitative behaviour of the dominant imaginary component, but remains quantitatively different from the non-perturbative one.

Because the amplitude of the third harmonic is small, the Monte Carlo estimation becomes significantly more demanding. For the dominant imaginary component, the non-perturbative Monte Carlo calculation reproduces the qualitative features of the response, namely the location of extrema near the interfaces and the sign changes

across them, but a non-negligible discrepancy with respect to the deterministic reference remains. Such a discrepancy is not explained by the estimated statistical uncertainty and its origin is presently unclear. At this stage, the comparison therefore supports only a qualitative agreement for the third harmonic, and further investigation would be required to clarify the source of the mismatch.




Figure 7.5: Third harmonic $\hat{\varphi}_3(x)$ of the vibrating rod model: real part (left) and imaginary part (right). Each panel shows the normalised harmonic (top) and the difference between Monte Carlo and deterministic solutions (bottom).

7.3.5 Summary of the comparison

Taken together, the vibrating rod results confirm the conclusions drawn from deterministic studies [13, 31]. For realistic vibration amplitudes, the linear noise formulation remains adequate for the first harmonic, but fails to describe the higher order harmonics correctly. The reformulated non-perturbative eigenvalue problem (4.17), combined with its Monte Carlo implementation in MGMC, yields results consistent with the deterministic non-perturbative reference for the dominant components of the first and second harmonics. For the third harmonic, the Monte Carlo results reproduce the qualitative interface-localised structure predicted by the deterministic non-perturbative solver, but a fully quantitative agreement is not achieved, and the remaining discrepancy is not accounted for by the estimated statistical uncertainty. This is both a strength and a limitation of the current implementation: it confirms the relevance of the non-perturbative formulation for vibration problems, while also indicating that the behaviour of higher-order harmonics still requires further investigation.

7.4 Parametric study

The previous section established the physical behaviour of the benchmark and the level of agreement obtained with the reformulated non-perturbative algorithm. The purpose of the present section is to clarify the numerical settings required for reliable non perturbative simulations, with particular attention to the choice of inactive generations and harmonic truncation. These two parameters play a central role in the stability and efficiency of the non-perturbative Monte Carlo method.

7.4.1 Inactive generations: monitoring k' and harmonic entropy

The non-perturbative Monte Carlo formulation is solved by power iteration in an enlarged phase space that includes the harmonic index. As in standard criticality calculations, a number of inactive generations must be discarded in order to allow the source to converge before tallies are accumulated. Source convergence is assessed by monitoring two diagnostics as a function of the generation index: the eigenvalue k' and an entropy indicator that measures how the particle population is distributed among harmonics. The latter is defined as the Shannon entropy along the harmonic coordinate,

$$H = - \sum_{h=0}^{h_{\max}} \frac{W_h}{W} \log_2 \left(\frac{W_h}{W} \right), \quad W = \sum_{h=0}^{h_{\max}} W_h, \quad (7.2)$$

where W_h denotes the total statistical weight carried by particles at harmonic h in a given generation. In practice, k' typically stabilises earlier than H , since the latter is sensitive to the slower redistribution of weight across the harmonic hierarchy.

Figure 7.6 reports these diagnostics for 20 independent replicas, shown as coloured curves, together with their ensemble average, shown in black. For the present configuration, H approaches a stationary regime after roughly 8×10^2 to 10^3 generations, whereas k' converges more rapidly. The required number of inactive generations is not universal and depends on the strength of the coupling between harmonics. This study confirms that the value $N_{\text{inactive}} = 15000$ used in Section 7.3 was sufficient for source convergence in the vibrating rod benchmark.




Figure 7.6: Convergence of the non-perturbative eigenvalue k' and the harmonic entropy H as functions of generation.

7.4.2 Harmonic truncation: choice of $h_{\max}$

The non-perturbative frequency-domain formulation yields, in principle, an infinite coupled system for $\{\hat{\varphi}_h\}_{h \geq 0}$, which must be truncated by selecting a maximum harmonic index $h_{\max}$. Increasing $h_{\max}$ reduces truncation bias, but at the same time spreads the finite particle population over a larger number of harmonics, thereby increasing the variance on the low order harmonics of interest. The choice of $h_{\max}$ is therefore a compromise between truncation error and statistical efficiency.

To quantify the residual truncation effect, the relative difference

$$\left| \frac{\Delta \hat{\varphi}_h}{\hat{\varphi}_h} \right| (x) = \frac{|\hat{\varphi}_h^{(h_{\max})}(x) - \hat{\varphi}_h^{(7)}(x)|}{|\hat{\varphi}_h^{(7)}(x)|}, \quad (7.3)$$

is considered, where $\hat{\varphi}_h^{(7)}$ denotes a reference non-perturbative calculation performed with $h_{\max} = 7$. Figures 7.7–7.9 report this quantity for $h = 1, 2, 3$ and for several values of $h_{\max}$. In all three figures, coloured solid lines correspond to non-perturbative simulations with different truncations ($h_{\max} = 1, 2, 3, 4$), the black dotted line shows the corresponding relative difference between linear and non-perturbative Monte Carlo solutions obtained with $h_{\max} = 7$, and the shaded band represents the estimated one-sigma Monte Carlo uncertainty on the relative difference estimator.




Figure 7.7: Effect of harmonic truncation on the non-perturbative solution for $h = 1$. Relative difference $|\Delta\hat{\varphi}_1/\hat{\varphi}_1|$ with respect to a reference non-perturbative calculation with $h_{\max} = 7$, as defined in Eq. (7.3).

For the first harmonic (Fig. 7.7), the relative differences remain small for all truncations considered. Even the smallest $h_{\max}$ values lead to deviations of at most order 10^{-2} over most of the domain. This confirms that the first harmonic is only weakly sensitive to the truncation strategy and that both the non-perturbative and linear formulations provide a consistent description of this component for the present benchmark.




Figure 7.8: Effect of harmonic truncation on the non-perturbative solution for $h = 2$. Relative difference $|\Delta\hat{\varphi}_2/\hat{\varphi}_2|$ with respect to a reference non-perturbative calculation with $h_{\max} = 7$, as defined in Eq. (7.3).

The behaviour of the second harmonic is shown in Fig. 7.8. For $h_{\max} < 3$, the

relative differences with respect to the $h_{\max} = 7$ reference remain noticeably larger than for $h_{\max} = 3$ and $h_{\max} = 4$, indicating that the coupling to higher harmonics is not yet fully resolved. Nevertheless, even these truncated non-perturbative solutions remain closer to the reference than the linear Monte Carlo result, whose discrepancy, shown by the black dotted line, remains the dominant error over most of the domain. By contrast, the curves for $h_{\max} = 3$ and $h_{\max} = 4$ lie comfortably below the linear versus non-perturbative discrepancy and within the one-sigma bands. Truncation effects are therefore subdominant with respect to the physical difference between the two formulations. This indicates that $h_{\max} \geq 3$ is required for a reliable estimate of $\hat{\varphi}_2$.




Figure 7.9: Effect of harmonic truncation on the non-perturbative solution for $h = 3$. Relative difference $|\Delta\hat{\varphi}_3/\hat{\varphi}_3|$ with respect to a reference non-perturbative calculation with $h_{\max} = 7$, as defined in Eq. (7.3).

For the third harmonic (Fig. 7.9), the relative difference profiles are strongly affected by the smallness of the signal and by the associated statistical uncertainty. The curves corresponding to $h_{\max} = 3$ and $h_{\max} = 4$ are essentially indistinguishable, and increasing $h_{\max}$ beyond 3 does not yield a clear and systematic reduction of $|\Delta\hat{\varphi}_3/\hat{\varphi}_3|(x)$ across the domain. This suggests that, for the present statistics, a truncation at $h_{\max} = 3$ already provides a consistent representation of the third harmonic.

In summary, the dedicated truncation study indicates that the harmonics $h = 0, 1, 2, 3$ discussed in Section 7.3 are satisfactorily represented once couplings up to $h_{\max} = 3$ are included. Accurate estimates of $h = 0$ and $h = 1$ are obtained already with small truncations, the second harmonic requires at least $h_{\max} = 3$ in order to resolve the relevant coupling structure, and the third harmonic is consistently represented for $h_{\max} \geq 3$.

7.4.3 Figure of Merit study

The computational performance of the non-perturbative and linear Monte Carlo calculations was assessed using the Figure of Merit introduced in Eq. (3.16) of Section 3.1.4. Since the quantities of interest in the vibrating rod benchmark are spatial profiles, the standard deviation entering the FoM was averaged over all detector positions. For the non-perturbative calculation, the statistical uncertainty was estimated from 20 independent replicas. Since the previous subsection established $h_{\max} = 3$ to be sufficient to represent the harmonics considered in this work, the non-perturbative simulations used for the FoM analysis were performed with this truncation. For the linear calculations, the statistical uncertainty was estimated from 10^4 pseudo-independent replicas, obtained by source re-sampling separated by 20 power iteration generations. A more conservative estimate was also considered by reducing the number of skipped generations in the linear calculations to 3. This choice slightly increased the variance but reduced the computation time by about 30%. The resulting FoM ratio, defined as $\text{FoM}_{\text{NP}}/\text{FoM}_{\text{lin}}$, together with the corresponding conservative estimates, are reported in Table 7.2.

Table 7.2: Figure of Merit ratios for the vibrating rod model.

	1 st harmonic	2 nd harmonic	Combined
Measured	2.9	1.2	4.0
Conservative estimate	2.0	0.8	2.8

The results show that, for the present benchmark and for the selected harmonics, the non-perturbative calculation is competitive not only in terms of accuracy, but also in terms of computational efficiency. For the first harmonic, using the non-perturbative formulation instead of the corresponding linear calculation improves the FoM by a factor of 2.9, or 2.0 according to the conservative estimate. For the second harmonic, the measured gain is smaller, with a FoM ratio of 1.2, while the conservative estimate gives 0.8. This latter comparison should be interpreted with care, since the linear formulation is not physically accurate for the second harmonic in the vibrating rod benchmark.

The main advantage appears when the first and second harmonics are considered together: the measured combined FoM ratio is 4.0, while the conservative estimate is 2.8. The non-perturbative calculation provides both harmonics simultaneously, whereas the linear formulation requires separate fixed-source calculations for each harmonic. As a result, obtaining the same precision on both harmonics with the linear approach would require a larger total computation time.

Chapter 8

Conclusions and Perspectives

This thesis investigated the non-perturbative Monte Carlo formulation of neutron noise in the frequency domain using the Monte Carlo mini-app MGMC, developed at CEA. To this end, systematic comparisons were performed between non-perturbative Monte Carlo simulations, the orthodox linear formulation, and deterministic reference solutions.

The homogeneous infinite-medium chapter first addressed a targeted verification of the flight-based scheme combined with the track-length estimator. A dedicated scattering-free high frequency test case was designed to make a previously identified implementation error in the FB+TL combination observable. Before correction, this combination showed statistically significant discrepancies with respect to the analytical reference, while the other scheme-estimator combinations remained compatible with it. After correction, the FB+TL results recovered agreement with the reference, confirming the corrected implementation of the track-length estimator in the flight-based scheme.

The same homogeneous framework was then used as a first simplified setting to validate the non-perturbative implementation and to compare it with the orthodox linear formulation. In the plateau region, the non-perturbative formulation generated non-zero higher harmonics and therefore a richer spectral response than the linear formulation, which remains restricted to the driving component. At high frequency, these higher harmonics were strongly damped and the two formulations became essentially indistinguishable within statistical uncertainty.

The most demanding benchmark considered in this thesis was the one-dimensional vibrating rod problem, representative of vibration-induced noise in systems with

moving interfaces. For this case, the original non-perturbative eigenvalue formulation exhibited poor convergence properties at realistic vibration amplitudes, so a reformulated non-perturbative eigenvalue problem was adopted and its Monte Carlo implementation was validated against deterministic reference solutions for the dominant components of the response. The results confirmed the conclusions suggested by deterministic studies: the zero frequency component and the first harmonic were reproduced accurately by both the linear and non-perturbative formulations, whereas the second harmonic clearly revealed the limitations of the orthodox formulation. The third harmonic provided a more demanding test for the non-perturbative Monte Carlo formulation: the main qualitative features of the deterministic non-perturbative response were recovered, but a fully quantitative agreement was not yet achieved at the current level of statistics.

The parametric study further clarified the numerical settings required for reliable non-perturbative simulations in the vibrating rod benchmark. Monitoring the reformulated eigenvalue k' together with the harmonic entropy proved effective for selecting a suitable number of inactive generations and for checking that the source had reached a stationary regime in harmonic space. The truncation analysis showed that, for the harmonics examined in this thesis, couplings up to $h_{\max} = 3$ were already sufficient to represent $h = 0, 1, 2, 3$ consistently.

A Figure of Merit analysis further showed that, with $h_{\max} = 3$, the non-perturbative calculation was computationally competitive with the corresponding linear calculations. The main advantage appeared when the first and second harmonics were considered together, since the non-perturbative formulation computes them simultaneously, leading to a FoM about four times larger than that of the linear calculation.

The study also highlights the current limitations of the method. Accurately resolving small higher-order harmonics remains difficult, as illustrated by the third harmonic in the vibrating rod benchmark. In addition, the weight cancellation strategy used in the present implementation remains closely tied to the simplified structure of the rod model, where the energy and angular variables are discrete and scattering is isotropic. Although this is sufficient for the benchmark considered here, extending the same strategy to more general transport configurations is expected to be significantly more difficult.

Several extensions naturally follow from the present work. On the algorithmic side, variance-reduction strategies specifically tailored to the non-perturbative formulation should be developed in order to improve the statistical quality of weak higher harmonics without excessive computational cost. Another important improvement would be the extension of weight cancellation to more general transport configurations, although this is expected to be considerably more difficult than in the simplified rod benchmark. The low-frequency regime should also be investigated

further, since the coupled harmonic system becomes increasingly ill-conditioned as the frequency decreases, making non-perturbative eigenvalue calculations more challenging. From the application point of view, the methodology could be tested on other simplified benchmarks and then extended progressively towards more realistic three-dimensional and continuous-energy reactor configurations.

In conclusion, this thesis shows that non-perturbative Monte Carlo neutron noise calculations in the frequency domain are feasible and relevant in regimes where harmonic coupling cannot be neglected. At the same time, the work shows that the accurate resolution of weak higher-order components is numerically demanding. The main perspective opened by this study is therefore not only to extend the range of applications of the method, but also to make non-perturbative Monte Carlo noise calculations sufficiently robust and efficient to become practical tools for more realistic reactor diagnostics.

Appendix A

Mathematical Background

A.1 Fourier transform

The Fourier transform provides a representation of a time dependent function in terms of its frequency content. In the present thesis it is used to move from the time domain to the frequency domain, where the neutron noise equations can be analysed harmonic by harmonic.

Let $f \in L^1(\mathbb{R})$. With the convention adopted throughout this manuscript, the Fourier transform of f is defined as

$$\hat{f}(\omega) = \mathcal{F}[f](\omega) = \int_{-\infty}^{+\infty} f(t) e^{-i\omega t} dt, \quad (\text{A.1})$$

where $\omega \in \mathbb{R}$ denotes the angular frequency. If $\hat{f} \in L^1(\mathbb{R})$, the inverse transform is

$$f(t) = \mathcal{F}^{-1}[\hat{f}](t) = \frac{1}{2\pi} \int_{-\infty}^{+\infty} \hat{f}(\omega) e^{i\omega t} d\omega. \quad (\text{A.2})$$

These definitions are consistent with the convention used in neutron noise theory in the main text and in the literature [34, 35, 12].

A.1.1 Remarks on function spaces

The definition in Eq. (A.1) is classical for $L^1(\mathbb{R})$ functions. The transform can also be extended to $L^2(\mathbb{R})$, where it is understood in the mean square sense. In that setting, the Plancherel identity states that

$$\int_{-\infty}^{+\infty} |f(t)|^2 dt = \frac{1}{2\pi} \int_{-\infty}^{+\infty} |\hat{f}(\omega)|^2 d\omega, \quad (\text{A.3})$$

which shows that the Fourier transform preserves the L^2 norm up to the factor imposed by the chosen convention [35].

A.1.2 Basic properties

The Fourier transform satisfies several elementary properties that are repeatedly used in frequency domain transport and noise theory. Let $f, g \in L^1(\mathbb{R})$, and let $a, b \in \mathbb{C}$. Then:

i) *Linearity*:

$$\mathcal{F}[af + bg](\omega) = a\hat{f}(\omega) + b\hat{g}(\omega).$$

ii) *Time shift*:

$$\mathcal{F}[f(t - t_0)](\omega) = e^{-i\omega t_0} \hat{f}(\omega).$$

iii) *Frequency shift*:

$$\mathcal{F}[e^{i\omega_0 t} f(t)](\omega) = \hat{f}(\omega - \omega_0).$$

iv) *Scaling*:

$$\mathcal{F}[f(at)](\omega) = \frac{1}{|a|} \hat{f}\left(\frac{\omega}{a}\right), \quad a \neq 0.$$

v) *Differentiation*: if f is sufficiently regular and $f^{(n)} \in L^1(\mathbb{R})$, then

$$\mathcal{F}\left[\frac{d^n f}{dt^n}\right](\omega) = (i\omega)^n \hat{f}(\omega).$$

vi) *Convolution theorem*: if $f * g$ denotes the convolution of f and g , then

$$\mathcal{F}[f * g](\omega) = \hat{f}(\omega) \hat{g}(\omega).$$

Two identities used repeatedly in the main text are

$$\mathcal{F}[\delta(t - t_0)](\omega) = e^{-i\omega t_0}, \tag{A.4}$$

$$\mathcal{F}[e^{i\omega_0 t}](\omega) = 2\pi \delta(\omega - \omega_0). \tag{A.5}$$

In particular,

$$\mathcal{F}^{-1}[1](t) = \delta(t), \tag{A.6}$$

with the present Fourier convention [34, 12].

A.1.3 Fourier transform of a periodic function

Let $f(t)$ be a periodic function of period $T > 0$, namely

$$f(t + T) = f(t), \quad \forall t \in \mathbb{R}, \quad (\text{A.7})$$

and let

$$\omega_0 = \frac{2\pi}{T} \quad (\text{A.8})$$

be the corresponding fundamental angular frequency. In general, a non zero periodic function does not belong to $L^1(\mathbb{R})$, so its Fourier transform is not defined in the classical sense. Its spectral content is instead represented through its Fourier series,

$$f(t) = \sum_{n=-\infty}^{+\infty} c_n e^{in\omega_0 t}, \quad (\text{A.9})$$

where the Fourier coefficients are

$$c_n = \frac{1}{T} \int_0^T f(t) e^{-in\omega_0 t} dt. \quad (\text{A.10})$$

Using Eq. (A.5), the Fourier transform of a periodic function is obtained in the distributional sense as

$$\hat{f}(\omega) = 2\pi \sum_{n=-\infty}^{+\infty} c_n \delta(\omega - n\omega_0). \quad (\text{A.11})$$

Hence a periodic signal has a discrete spectrum concentrated at the harmonics $n\omega_0$ of the fundamental frequency [34, 12]. This representation is the natural starting point for the harmonic formulations used in frequency domain neutron noise theory.

A.2 Convolution product

A.2.1 Continuous convolution

Let $f, g \in L^1(\mathbb{R})$. Their convolution is defined by

$$(f * g)(t) = \int_{-\infty}^{+\infty} f(\tau) g(t - \tau) d\tau. \quad (\text{A.12})$$

When the integral exists, $f * g$ is again a function on $\mathbb{R}$. Convolution is commutative, associative, and distributive over addition:

$$f * g = g * f, \quad (\text{A.13})$$

$$f * (g * h) = (f * g) * h, \quad (\text{A.14})$$

$$f * (g + h) = f * g + f * h. \quad (\text{A.15})$$

The importance of convolution in Fourier analysis follows from the convolution theorem already stated in Section A.1.2, namely

$$\mathcal{F}[f * g](\omega) = \hat{f}(\omega)\hat{g}(\omega). \quad (\text{A.16})$$

Conversely, the Fourier transform of a product is

$$\mathcal{F}[fg](\omega) = \frac{1}{2\pi}(\hat{f} * \hat{g})(\omega), \quad (\text{A.17})$$

which explains why products in the time domain become convolutions in the frequency domain [34, 35, 12].

A.2.2 Discrete convolution

Let $u, v : \mathbb{Z} \rightarrow \mathbb{C}$. Their discrete convolution is defined as

$$(u * v)[n] = \sum_{m=-\infty}^{+\infty} u[m] v[n - m]. \quad (\text{A.18})$$

If the sequences have finite support, the sum reduces to a finite number of terms.

Discrete convolution is the natural counterpart of Eq. (A.12) when the spectrum is represented by a discrete set of harmonics. In particular, when a periodic perturbation is expanded over the frequencies $n\omega_0$, the continuous frequency convolution reduces to a discrete coupling among harmonic indices. This is precisely the structure that appears in the non perturbative neutron noise equations discussed in the main text.

Bibliography

- [1] Imre Pázsit and Christophe Demazière. «Noise Techniques in Nuclear Systems». In: *Handbook of Nuclear Engineering*. Ed. by Dan Gabriel Cacuci. Vol. 3. Boston, MA: Springer, 2010. Chap. 14. DOI: 10.1007/978-0-387-98149-9_14 (cit. on pp. 1, 18, 19, 31–33, 37, 38).
- [2] Joseph A. Thie. *Power Reactor Noise*. La Grange Park, IL: American Nuclear Society, 1981 (cit. on p. 1).
- [3] M. M. R. Williams. *Random Processes in Nuclear Reactors*. Oxford: Pergamon Press, 1974 (cit. on p. 1).
- [4] Christophe Demazière. «CORE SIM: A Multi-Purpose Neutronic Tool for Research and Education». In: *Annals of Nuclear Energy* 38.12 (2011). DOI: 10.1016/j.anucene.2011.06.010 (cit. on p. 1).
- [5] Amélie Rouchon. «Analyse et développement d’outils numériques déterministes et stochastiques résolvant les équations du bruit neutronique». PhD Thesis. Université Paris-Saclay, 2016 (cit. on pp. 1, 2, 35, 53).
- [6] Andrea Gammicchia, Simone Santandrea, Igor Zmijarevic, Richard Sanchez, Zoran Stankovski, Sandra Dulla, and Paolo Mosca. «A MOC-Based Neutron Kinetics Model for Noise Analysis». In: *Annals of Nuclear Energy* 137 (2020). DOI: 10.1016/j.anucene.2019.107070 (cit. on p. 1).
- [7] Paul Cosgrove, Maximilian Kraus, and Valeria Raffuzzi. «A Memory-Efficient Neutron Noise Algorithm for Reactor Physics». In: *Annals of Nuclear Energy* 201 (2024). DOI: 10.1016/j.anucene.2024.110450 (cit. on p. 1).
- [8] T. Yamamoto. «Monte Carlo Method with Complex-Valued Weights for Frequency Domain Analyses of Neutron Noise». In: *Annals of Nuclear Energy* 58 (2013) (cit. on pp. 1, 41, 48).
- [9] Amélie Rouchon, Andrea Zoia, and Richard Sanchez. «A New Monte Carlo Method for Neutron Noise Calculations in the Frequency Domain». In: *Annals of Nuclear Energy* 102 (2017). DOI: 10.1016/j.anucene.2016.11.035 (cit. on pp. 1, 40, 41, 48).

- [10] Hunter Belanger, Davide Mancusi, Amélie Rouchon, and Andrea Zoia. «Variance Reduction and Noise Source Sampling Techniques for Monte Carlo Simulations of Neutron Noise Induced by Mechanical Vibrations». In: *Nuclear Science and Engineering* 197.4 (2023) (cit. on pp. 1, 45).
- [11] Imre Pázsit. «Investigation of the Space-Dependent Noise Induced by a Vibrating Absorber». In: *Atomkernenergie* 30 (1977) (cit. on p. 1).
- [12] Richard Sanchez. «Some Comments in Neutron Noise Theory». In: *Annals of Nuclear Energy* 86 (2015). DOI: 10.1016/j.anucene.2015.04.003 (cit. on pp. 1, 15, 32, 34, 81–84).
- [13] Andrea Zoia, Amélie Rouchon, Baptiste Gasse, Christophe Demazière, and Paolo Vinai. «Analysis of the Neutron Noise Induced by Fuel Assembly Vibrations». In: *Annals of Nuclear Energy* 154 (2021). DOI: 10.1016/j.anucene.2020.108061 (cit. on pp. 1, 2, 15, 33–35, 63, 64, 68, 70).
- [14] Axel Fauvel, Amélie Rouchon, Davide Mancusi, and Andrea Zoia. «Monte Carlo Sampling Strategy for a Non-Perturbative Approach to the Neutron Noise Equation: A Proof of Concept». In: *Proceedings of the International Conference on Mathematics and Computational Methods Applied to Nuclear Science and Engineering (M&C 2025)*. Denver, CO, USA: American Nuclear Society, Apr. 2025 (cit. on pp. 2, 35, 36, 40, 43–45, 53).
- [15] James J. Duderstadt and Louis J. Hamilton. *Nuclear Reactor Analysis*. New York: John Wiley & Sons, 1976 (cit. on pp. 5–10, 13, 16, 17).
- [16] John R. Lamarsh and Anthony J. Baratta. *Introduction to Nuclear Engineering*. 3rd. Upper Saddle River, NJ: Prentice Hall, 2001 (cit. on pp. 5–9, 11, 12).
- [17] Kenneth S. Krane. *Introductory Nuclear Physics*. New York: John Wiley & Sons, 1987. ISBN: 978-0-471-80553-3 (cit. on pp. 5–7).
- [18] Neil E. Todreas and Mujid S. Kazimi. *Nuclear Systems I: Thermal Hydraulic Fundamentals*. 1st. New York: Hemisphere Publishing Corporation, 1990. ISBN: 0-89116-935-0 (cit. on p. 6).
- [19] Bertrand Barré. «Gas-Cooled Reactors». In: *Handbook of Nuclear Engineering*. Ed. by Dan Gabriel Cacuci. Vol. 4. Boston, MA: Springer, 2010. Chap. 22. DOI: 10.1007/978-0-387-98149-9_22 (cit. on p. 6).
- [20] Weston M. Stacey. *Nuclear Reactor Physics*. Second Edition, Completely Revised and Enlarged. Weinheim, Germany: Wiley-VCH, 2007. ISBN: 978-3-527-40679-1 (cit. on pp. 6, 7, 12, 18, 19, 25, 27).
- [21] George I. Bell and Samuel Glasstone. *Nuclear Reactor Theory*. New York: Van Nostrand Reinhold, 1970 (cit. on pp. 8, 11, 14).

- [22] Anil K. Prinja and Edward W. Larsen. «General Principles of Neutron Transport». In: *Handbook of Nuclear Engineering*. Ed. by Dan Gabriel Cacuci. Vol. 1. Boston, MA: Springer, 2010. Chap. 5. DOI: 10.1007/978-0-387-98149-9_5 (cit. on pp. 8, 10–14, 17, 21).
- [23] Robert C. Block, Yaron Danon, Frank Gunsing, and Robert C. Haight. «Neutron Cross Section Measurements». In: *Handbook of Nuclear Engineering*. Ed. by Dan Gabriel Cacuci. Vol. 1. Boston, MA: Springer, 2010. Chap. 1. DOI: 10.1007/978-0-387-98149-9_1 (cit. on pp. 9, 11).
- [24] Alireza Haghighat. *Monte Carlo Methods for Particle Transport*. Boca Raton, FL: CRC Press, 2015. ISBN: 978-1-4665-9253-7 (cit. on pp. 21, 23–28).
- [25] Jerome Spanier and Ely M. Gelbard. *Monte Carlo Principles and Neutron Transport Problems*. Unabridged republication of the 1969 edition published by Addison-Wesley. Mineola, New York: Dover Publications, 2008. ISBN: 9780486462936 (cit. on pp. 21–23, 25–28).
- [26] Athanasios Papoulis. *Probability, Random Variables, and Stochastic Processes*. 3rd. New York: McGraw-Hill, 1991. ISBN: 978-0-07-048477-1 (cit. on pp. 22, 23, 25).
- [27] John F. Kenney. *Mathematics of Statistics, Part One*. New York: D. Van Nostrand Company Inc, 1939 (cit. on p. 24).
- [28] Ivan Lux and Laszlo Koblinger. *Monte Carlo Particle Transport Methods: Neutron and Photon Calculations*. Boca Raton, FL: CRC Press, 1991. ISBN: 0-8493-6074-9 (cit. on pp. 25, 28, 29).
- [29] Hunter Belanger, Davide Mancusi, and Andrea Zoia. «Exact Weight Cancellation in Monte Carlo Eigenvalue Transport Problems». In: *Physical Review E* 104.1 (2021), p. 015306. DOI: 10.1103/PhysRevE.104.015306 (cit. on p. 45).
- [30] Larry Wasserman. *All of Statistics: A Concise Course in Statistical Inference*. New York: Springer, 2004. ISBN: 978-0-387-40272-7 (cit. on p. 49).
- [31] Amélie Rouchon and Richard Sanchez. «Analysis of Vibration-Induced Neutron Noise Using One-Dimension Noise Diffusion Theory». In: *Proceedings of the International Congress on Advances in Nuclear Power Plants (ICAPP 2015)*. Paper 15098. Nice, France, 2015 (cit. on pp. 63, 68, 70).
- [32] Youcef Saad and Martin H. Schultz. «GMRES: A Generalized Minimal Residual Algorithm for Solving Nonsymmetric Linear Systems». In: *SIAM Journal on Scientific and Statistical Computing* 7.3 (1986). DOI: 10.1137/0907058 (cit. on p. 66).

- [33] Richard B. Lehoucq, Danny C. Sorensen, and Chao Yang. *ARPACK Users' Guide: Solution of Large-Scale Eigenvalue Problems with Implicitly Restarted Arnoldi Methods*. SIAM, 1998. DOI: 10.1137/1.9780898719628 (cit. on p. 66).
- [34] T. W. Körner. *Fourier Analysis*. Cambridge: Cambridge University Press, 1988. ISBN: 978-0-521-25120-4 (cit. on pp. 81–84).
- [35] Walter Rudin. *Real and Complex Analysis*. 3rd. New York: McGraw-Hill, 1987. ISBN: 9780070542341 (cit. on pp. 81, 82, 84).